

Data protection with Veeam Backup & Replication, HPE StoreOnce, and HPE StoreEver



Contents

Executive summary.....	3
Solution overview.....	4
Solution components.....	5
Hardware.....	5
Software	6
Best practices and configuration guidance for the solution.....	6
Introduction to HPE StoreOnce security concepts.....	7
Enabling Dual Authorization for HPE StoreOnce.....	7
Creating a Catalyst Store to be used as a Veeam backup repository	8
Catalyst Store Advanced Settings.....	10
Configuring access permissions to an HPE StoreOnce Catalyst Store	11
How to connect to an HPE StoreOnce Catalyst Store via IP or Fibre Channel.....	11
Initial Fibre Channel configuration.....	12
Deploying a Veeam gateway server.....	13
Configuring the Veeam gateway server for CoFC.....	15
Creating an immutable HPE StoreOnce Catalyst-based Veeam backup repository.....	16
Creating multiple Veeam backup repositories inside the same HPE StoreOnce Catalyst Store	24
Creating a primary backup job for VMs.....	26
Creating a Veeam backup job for Windows and Linux computers.....	33
Validation of backup immutability.....	36
Copying backups to an off-site HPE StoreOnce and Cloud Bank Storage.....	39
Recovering data from an HPE StoreOnce Catalyst repository	52
Migrating Veeam backup repositories to/from an HPE StoreOnce Catalyst Store	52
Configuring Veeam Tape Jobs with HPE StoreEver	60
Recovering data from HPE StoreEver tape media	67
Performance and tuning considerations.....	68
Performance examples.....	68
Summary.....	69
Resources, contacts, or additional links	71



Executive summary

Modern businesses are anchored to a hybrid, hyper-scale ecosystem where every interaction results in massive and complex datasets that are crucial in their digital transformation journey. While this data needs to be protected and recovered in a simple, automated, reliable, and resilient manner, an unprecedented surge in ransomware attacks mandates the inclusion of data protection within a comprehensive security strategy.

Hewlett Packard Enterprise and Veeam co-innovation provides numerous integration options with Veeam Backup & Replication and HPE storage and server portfolios to deliver features that enable rapid, efficient, scalable, and secure backup and recovery throughout the lifecycle of data generated by a myriad of workloads. This technical paper demonstrates the best practices and benefits of integrating [HPE StoreOnce](#) and [HPE StoreEver](#) with Veeam Backup & Replication version 12 for complying with the 3-2-1-1 data protection rule (extension of [3-2-1 rule](#), where one copy must be immutable or air-gapped for ransomware protection) to protect data against planned and unplanned outages.

Benefits of the integration include:

- **Reliable and fast data protection with reduced backup data footprint:** HPE StoreOnce systems integrate with Veeam Backup & Replication via the Catalyst protocol, which provides unmatched source-side deduplication capabilities. This leads to a more efficient backup of deduplicated data transfer over LAN, WAN, and Fibre Channel (FC) connections with substantial storage space savings (~95%). The new fixed-block chunking for pre-aligned Veeam backup data improves single-stream backup by over 300%. The underlying flash tier of HPE StoreOnce systems further expedites synthetic full backup creation and Instant/granular recovery operations.
- **Fast and reliable backup replication:** Veeam-managed storage-based replication across any number of HPE StoreOnce units efficiently offloads all data transfer operations to HPE StoreOnce Catalyst Copy. Due to the deep integration with Veeam Backup & Replication software, the backup replication is granularly managed by Veeam, and all backup copies are immediately accessible on the Veeam console. Storage-based replication is faster, and it also reduces costs as it requires less network bandwidth. Finally, Catalyst Copy is executed by HPE StoreOnce and does not require additional components such as Veeam Gateway or WAN acceleration servers.
- **Affordable and flexible archival of enterprise data in the cloud:** HPE StoreOnce Cloud Bank Storage (CBS) provides a means for creating off-site copies of data in multiple public or on-premises low-cost cloud object storage targets for Disaster Recovery (DR)/Long Term Retention (LTR) without relying on an appliance or a storage gateway. When a CBS is configured as the secondary target for a Veeam-managed Catalyst Copy job, HPE StoreOnce performs a storage-based replication of highly deduplicated backup data to public and private cloud object stores. This results in a faster process because less data is transferred, and up to a 5x cost reduction on the cloud storage monthly fee.
- **Robust ransomware protection:** Data written into HPE StoreOnce Catalyst Stores are invisible to host systems and are unalterable through retention lock and object-level immutability set by Veeam primary backup jobs leveraging the Catalyst client APIs. This coupled with HPE StoreOnce dual authorization and compliance mode provides fool-proof, end-to-end protection against cryptographic and key-logger attacks intended for modification and deletion of valuable backup data.
- **Low cost, scalable, air-gapped Long Term Retention(LTR) with assured longevity:** The [HPE StoreEver](#) tape library family provides a means to the inexpensive, scalable, and highly dense archival of cold data in a sustainable manner with a shelf life of over 30 years. HPE StoreEver offers a broad portfolio from the entry level HPE MSL 1/8 autoloader; up to the HPE MSL6480 that is capable of archiving up to 25.2 PB with LTO-9 read/write or WORM tape media. The tape media can be automatically vaulted and write-protected by Veeam Tape Jobs, thus qualifying them to be perfect targets for long-term air-gapped data retention.

Target audience: Presales consultants, solution architects, and backup administrators who are designing and implementing a holistic data protection solution with HPE StoreOnce, HPE StoreEver, and Veeam Backup & Replication software. Working knowledge of Veeam Backup & Replication, HPE StoreOnce, and HPE StoreEver is recommended.

Document purpose: This paper describes a solution that includes best practices and highlights recognizable benefits to technical audiences. It provides information for setting up and managing data availability infrastructures that include Veeam Backup & Replication and HPE StoreOnce backup appliances configured with HPE StoreOnce Catalyst backup targets. This document is current to Veeam Backup & Replication version 12 and HPE StoreOnce firmware version 4.3.2. For more details about Veeam Backup & Replication, see the [Veeam Backup & Replication User Guide](#). For HPE StoreOnce information, see hpe.com/storage/StoreOnce. For more information on HPE StoreEver MSL libraries, see hpe.com/storage/StoreEver.

This paper describes solution testing performed by Hewlett Packard Enterprise in Sept 2022.



Solution overview

This solution focuses on the best practices and benefits of deploying the HPE StoreOnce systems as primary and secondary backup targets with Veeam Backup & Replication version 12 for creating both on-site and off-site backup copies. The solution was designed keeping the 3-2-1-1 rule (three copies of data on two different media—one offsite and one offline/immutable) in mind. This paper also provides best practices and performance metrics for additional backup copies to HPE StoreEver tape libraries for offline LTR. Use cases covered in this document include:

- Creation of immutable backups with HPE StoreOnce as the primary backup target
- Scheduled creation of secondary backup copies in a remote/off-site HPE StoreOnce backup appliance using Veeam-managed Catalyst Copy Jobs (CCJ)
- Configuring Veeam-managed CCJs to create backup copies in reliable and cost-efficient HPE StoreOnce Cloud Bank object store for DR and LTR.
- Use of HPE StoreOnce as the secondary backup target for a Veeam Backup Copy Job (BCJ)
- Automated backup and archival process using Veeam Tape Jobs and HPE StoreEver
- Migrating traditional disk-based repositories to highly efficient and deduplicated HPE StoreOnce Catalyst repositories

The solution environment layout in [Figure 1](#) illustrates the Veeam Backup & Replication architecture with data movement as follows:

1. The Veeam backup server sends a request to the proxy server to quiesce VMs and create a VM snapshot.
2. [Optional] Veeam integration with HPE primary storage ([HPE Alletra](#), [HPE Primera](#), [HPE Nimble](#), and [HPE 3PAR](#)) that hosts the virtualized or physical production infrastructure. This allows the Veeam backup server to directly communicate with the arrays to create application/crash-consistent snapshots that act as a data source for storage snapshot-integrated backup jobs. The VM snapshot is immediately deleted post storage snapshot creation.

Hardware snapshot integration is a Veeam best practice to ensure customers can experience full-featured restore from storage snapshots in addition to realizing greater backup performance. Primary snapshot Integration is highly recommended as it reduces the stun-effect on the production VMs and improves the backup performance by approximately 30%. Also, when production infrastructure is hosted on HPE primary storage, you can configure snapshot-only jobs with Veeam Backup & Replication to create application consistent storage snapshots that help adhere to stringent RPO and RTO requirements. The majority of the Veeam's advanced restore features, such as Instant VM Recovery (IVMR) using hardware snapshots, enable the recovered VMs to run at production speed. For additional information on data protection with hardware snapshot orchestration with HPE primary storage, see the [HPE Reference Configuration for Veeam Availability Suite with HPE Primera Storage](#) and the [HPE Reference Configuration for Veeam Availability Suite with HPE Nimble Storage](#).

Note

Without backups from storage snapshots, Veeam strongly recommends a two-tier solution for backup infrastructure. Backup to a regular disk-based repository like HPE Nimble or [HPE Apollo](#) storage is recommended with backup copies written to the HPE StoreOnce backup appliance for long-term backup data archival. In a two-tier approach, Veeam could target HPE Nimble or HPE Apollo storage as the primary backup target before using a Backup Copy Job to copy the backup to the HPE StoreOnce appliance. For customers wanting full featured restore options from backup files (in addition to storage snapshots), this two-tier solution is a Veeam best practice. See this [Veeam Knowledge Base](https://www.veeam.com/kb1745) article for more direction on the use of deduplication appliances in Veeam Backup Infrastructure: [veeam.com/kb1745](https://www.veeam.com/kb1745).

3. The Veeam backup server sends a request to the proxy server to read and transform the data to be backed up. Data is read either from attached storage snapshots or through the virtualization hypervisor if storage snapshot integration is not available in a backup job.
4. The gateway server hosting the HPE StoreOnce Catalyst client and Veeam data-mover service performs source-side deduplication of the data transformed by the proxy server before backing it up in a primary backup target configured using a local HPE StoreOnce Catalyst Store via an Ethernet or FC connection.
5. The Veeam backup server triggers any of the following user-defined auxiliary jobs for creating offline backup copies:
 - a. Veeam-managed HPE StoreOnce Catalyst Copy job—a Catalyst protocol-based storage replication job for copying data to an off-site HPE StoreOnce repository
 - b. Veeam-managed HPE StoreOnce Catalyst Copy job—a Catalyst protocol-based storage replication job for copying data to HPE StoreOnce Cloud Bank Storage



6. The Veeam Tape Server reads data from the source HPE StoreOnce backup repository via the gateway server and sends it to HPE StoreEver tape storage.




Figure 1. Veeam Backup & Replication data flow to HPE storage devices

In the case of HPE StoreOnce, Veeam does not install a data mover on the HPE StoreOnce appliance, and instead uses a gateway service to run the Catalyst protocol. If no additional Windows servers are added, Veeam will use the default backup proxy for this role. A Windows Agent can also run the gateway component that executes HPE StoreOnce Catalyst, which can be desirable for large, highly transactional servers. The same HPE StoreOnce Catalyst Store can be accessed by multiple gateway services concurrently (a new feature in Veeam Backup & Replication version 12) and Veeam can dynamically distribute the workload across multiple servers running the proxy and gateway components.

Figure 1 shows a simplified functional diagram of the architecture. In a production environment there could be multiple physical and virtual servers running Veeam proxy and gateway services, and the communication between the proxy and the gateway components generally occurs inside the same server without an additional hop in LAN.

Solution components

This section describes the hardware and software components used to create this solution.

Hardware

VMware ESXi Server

An [HPE ProLiant](#) DL380 Gen10 server with 32 cores and 512 GB RAM was used for installation of VMware ESXi™, 7.0.2. LAN connectivity was provided through two Emulex Corporation HP NC552SFP Dual Port 25GbE adapters. The ESXi server was used to host Veeam 12 Backup Server, proxy (for VM backups in hot-add mode), and gateway server VMs.

HPE ProLiant DL380 servers

Two identical HPE ProLiant DL380 Gen10 servers with 18 cores and 256 GB RAM servers were also used to host the Veeam proxy and gateway services to test snapshot-integrated backup jobs. The servers were connected to HPE Alletra 9000 storage through 16 Gb FC connectivity to directly read the backup data from the storage array snapshots for SAN backups.

HPE StoreOnce 5260

For this solution, an [HPE Gen4+ HPE StoreOnce Purpose Built Backup Appliance 5260](#) with 25GbE connectivity to the Veeam proxy and gateway servers was used. HPE Gen4+ hardware appliances with refreshed hardware and flash tier for persistent storage of metadata significantly reduce the latency and boost the performance during IVMR, granular restore, and synthetic full backup operations.

The HPE StoreOnce Gen4+ appliances are configured with [HPE StoreOnce firmware version 4.3.2](#), the minimum requirement for support of HPE StoreOnce Cloud Bank Storage, immutable Veeam Catalyst repositories, dual authorization, and fixed-block chunking (4 kb block size) for pre-aligned backup data by Veeam Backup & Replication Version 12. The firmware also comes with a highly optimized deduplication engine that improves single Catalyst Store backup performance and significantly reduces latency during random I/O and house-keeping operations.

Tape Server

An HPE ProLiant DL380 Gen10 server with 18 cores and 256 GB RAM was used for configuring the Tape Server with Windows Server 2019. LAN connectivity was provided through two HPE Ethernet 10/25Gb 2-port 640SFP28 adapters.

HPE Alletra 9000 storage

Storage capacity was provided by an [HPE Alletra 9000 array](#) that was interconnected via a 16 Gb FC fabric. VMware® datastores used storage from the HPE Alletra 9000 volumes to host the VMs that are backed up by Veeam Backup & Replication software. [HPE Nimble Storage](#) and [HPE Primera](#) arrays can also be used.

HPE StoreEver MSL3040 tape library

The HPE StoreEver MSL3040 tape library is an economical, scalable, and easily upgradable tape storage unit with a capacity of 28.8 PB with LTO-9. The MSL tape library was used for air-gapped LTR for primary backup target data. The tape library had firmware version 3290 and the LTO-9 tape drive had firmware version P371 and was interfaced to the Veeam Tape Server through 16 GB FC connectivity.

Software

HPE StoreOnce VSA

The functionalities supported were also validated with an [HPE StoreOnce VSA](#) that was deployed on a VMware ESXi 7.0.2 server and configured with [software version 4.3.2](#).

For more information, see [HPE StoreOnce system software](#)

Veeam Backup & Replication

Veeam Backup & Replication version 12 was used in this solution. All the Veeam software components—backup server, proxy servers, gateway servers, and tape server—were installed on the Windows 2019 OS.

For information on platform support, see [Veeam Backup & Replication 12 - Platform Support](#).

HPE StoreEver Management Software

HPE StoreEver Management Software provides an easy-to-use interface for efficiently managing, monitoring, and configuring an entire tape library environment. The HPE StoreEver Management Software solution includes HPE Command View for Tape Libraries (CVTL), HPE StoreEver TapeAssure Advanced, and HPE StoreEver Data Verification. See [HPE StoreEver Management Software](#) for more information about HPE StoreEver Management Software.

Best practices and configuration guidance for the solution

This section provides best practices and guidance to build the solution environment. Tuning parameters are also provided in some instances.

Hewlett Packard Enterprise recommends the following links for proper setup of the software and security components to adequately run Veeam Backup & Replication with HPE StoreOnce and HPE StoreEver:

- [HPE StoreOnce Support Matrix](#)
- [HPE StoreEver Support Matrix](#)
- [Veeam Backup & Replication - VMware vSphere platform support](#)
- [Veeam Backup & Replication - HPE StoreOnce appliance as a backup repository](#)
- [Veeam Backup & Replication - System Requirements for Veeam Agents](#)



Introduction to HPE StoreOnce security concepts

HPE StoreOnce has multiple advanced protection mechanisms to prevent data from being encrypted, altered, or deleted by a cyber or ransomware attack.

- **Catalyst data communication protocol:** Catalyst is the data communication protocol used by Veeam Backup & Replication to transfer data to HPE StoreOnce and to control additional functionalities, such as deduplication, replication, and immutability. HPE StoreOnce also supports NFS and SMB file-sharing protocols and it can emulate tape libraries (VTL). But it is recommended to not use these protocols with Veeam Backup & Replication as the Catalyst protocol offers more performance and optimization for Veeam workload, advanced integrated features, source-side deduplication, and immutability.
- **Immutability:** Immutable data is data that cannot be modified or deleted by any client application, backup administrators, or malicious attackers who have access to the data. Starting from HPE StoreOnce firmware version 4.3.2 and Veeam Backup & Replication version 12, Veeam can set the immutability period to back up files created in HPE StoreOnce Catalyst Store using the Catalyst APIs. An immutable backup file/object cannot be deleted or altered until the expiration of the immutable timestamp. This makes the backups on HPE StoreOnce Catalyst Stores tamper-proof against unauthorized or accidental deletions.
- **Compliance:** A storage device that archives data marked as immutable should protect the data not only from client applications and hosts with storage access via every data access protocol such as Catalyst, but it should also protect data from storage administrators who use the storage console to access the storage array for deleting volumes and files or for removing the immutability attributes. Starting from HPE StoreOnce firmware version 4.3.2, there is a “dual authorization” mechanism that makes it impossible for any HPE StoreOnce administrator to delete immutable data volumes (also known as Catalyst Stores). After the dual authorization is enabled, every HPE StoreOnce console/administrative operation that can lead to the deletion of a Catalyst Store—or any other system disruptive operation—requires a double authorization to be executed. Note that there is no way to delete single backup files or objects. More details are provided in the subsequent sections.
- **No root accounts on HPE StoreOnce:** Hewlett Packard Enterprise never discloses the HPE StoreOnce root account. Root credentials are never provided to the end user. The end user never needs to use the root account for any activity. Nobody can use the root account to break the immutability and compliance security measures.
- **No shell:** HPE StoreOnce does not provide shell access. Users cannot use known or unknown shell vulnerabilities to make privilege escalation.

Note

Starting with HPE StoreOnce generation 4 systems (October 2018), Catalyst is pre-enabled on all systems and does not require any additional HPE StoreOnce licensing.

Enabling Dual Authorization for HPE StoreOnce

HPE StoreOnce supports dual authorization from firmware version 4.3.2. Dual authorization is a system-level setting that requires the approval of a security officer to authorize specific disruptive operations, such as deleting/modifying a store, modifying system settings, and so on. If subscription is enabled, the security officer also receives email notifications intimating the necessity to approve or deny an action. This feature protects against data loss caused by accidental actions or ransomware attacks. It requires the credentials of two users: the Admin user who initiates the action and the Security Officer who approves/denies the action to perform critical actions. Thus, the probability of backup destruction is highly reduced.

To enable dual authorization at the HPE StoreOnce appliance level:

1. Log in to the HPE StoreOnce Management Console with administrator privileges.
2. Navigate to **Dual Authorization** → **Configure Dual Authorization** and slide the toggle bar to enable Dual Authorization as shown in Figure 2.






Figure 2. Enabling Dual Authorization for HPE StoreOnce appliance

Note

Enabling dual authorization at the HPE StoreOnce system level is a mandatory step before creating an immutable Catalyst-based backup repository in Veeam Backup & Replication version 12. Hewlett Packard Enterprise recommends the creation of more than one security officer for approving/denying dual authorization requests. The security officer credentials can be printed on paper and kept securely in closed envelopes. When a security officer is required, one of the envelopes can be opened and the security officer credentials can be used for approving/denying an operation queued-up for authorization. After this use, it is recommended to delete the account. This approach provides a stronger defence against malicious attacks.

Creating a Catalyst Store to be used as a Veeam backup repository

Before configuring a Veeam backup repository, an HPE StoreOnce Catalyst Store must be created on an HPE StoreOnce appliance that will store the backups. To do this:

1. Log in to the HPE StoreOnce Management Console (HPE StoreOnce GUI) with administrator privileges.
2. Navigate to **Data Services** → **Catalyst Stores** → **Create Store**.
3. Enter the **Catalyst Store Name** and click the edit button (pen) for **Security Settings** and set **Maximum ISV Controlled Data Immutability Retention** as shown in [Figure 3](#).
4. Do not enable **Server Controlled Data Immutability**. This feature is used by other ISVs not supporting the ISV Controlled Data Immutability and makes all the objects in the entire store immutable with a uniform retention period as defined by the administrator.
5. Do not enable **Store Encryption** as shown in [Figure 3](#) when not required as it adds unnecessary load to the system. If data-in-rest encryption is required, enable **Store Encryption** by editing the **Security Settings** of the Catalyst Store. HPE StoreOnce encryption is a free license included on all HPE StoreOnce Gen 4 models and newer. The license is available in most countries, except for specific countries where the license is not available because of export restrictions. While editing the **Security Settings** of a pre-created store, do not enable **Secure Erase**. Secure Erase on Catalyst Stores will have an impact on the performance due to increased disk I/O. Enable it only when needed and disable it when it is not needed.

Note

Updating the ISV-controlled immutability duration will not affect the existing Catalyst items. After the expiry of the ISV set immutable time duration, the Catalyst objects are retained in a read-only state. It is mandatory to use ISV-controlled Immutability with dual authorization as any attempt to modify the immutable epoch timestamp against any Veeam-created Catalyst backup objects (VM backups, server agent backups, database agent backups—RMAN, SAP®, SAP HANA®, Oracle®, MS SQL backups) created at the HPE StoreOnce level gets effectively blocked by dual authorization.






Figure 3. New Catalyst Store security configuration settings in the HPE StoreOnce GUI

6. If necessary, change the store **Advanced Settings** as shown in Figure 4, otherwise, keep the default setting. Make sure that the **Enforce Fixed Block Chunking** option is enabled for higher performance while backing up pre-aligned backup data. Click **Create** to complete the store creation.




Figure 4. HPE StoreOnce GUI showing a Catalyst Store configured with advanced settings



Catalyst Store Advanced Settings

The following settings are recommended in the advanced settings for better backup and restore performance:

- **Source-side deduplication (Low Bandwidth)** as the **Primary Transfer Policy**: LBW transfer processes the backup data on the Veeam gateway server to identify and only transfer unique backup data to the HPE StoreOnce appliance. This is called “source-side deduplication.” Source-side deduplication significantly reduces the physical backup data transferred over the network. It can also reduce backup time by increasing throughput, particularly when there are multiple concurrent backup jobs. It should also be used if the backup or Backup Copy Job is over a WAN because of the significant reduction in bandwidth requirements.

Source-side and target-side deduplication (HBW) produce the same level of data reduction. Source-side deduplication (LBW) is preferred because it increases the overall backup throughput and reduces the network traffic. Alternatively, HBW, might offer better performance only with very uncommon workloads—with one or very few concurrent streams and when **Enforce Fixed-Block Chunking** is not enabled. LBW is recommended for most use cases and should be specified unless there is a verified reason for not using it.

The network load reduction produced by LBW on the first backup of a VM is lower than for subsequent backups. Starting with the second backup, the network load reduction will be measurable. For restore/read activities, Catalyst does not produce any bandwidth reduction and data is sent rehydrated to the Veeam server.

- Enable **Enforce Fixed Block Chunking**: This is an enhancement introduced in HPE StoreOnce firmware 4.3.2 specifically for Veeam Backup & Replication Version 12. Enabling this option at the HPE StoreOnce Catalyst Store level for pre-aligned backup data by Veeam improves the single-stream backup performance by 4 to 5 times. Fixed Block Chunking also significantly improves the synthetic full backup performance. This gives a huge advantage by reducing the backup time of very large VMs.

This option modifies the HPE StoreOnce Catalyst deduplication algorithm to use fixed-block segments of length 4 kb for deduplication rather than using the default variable-segment-length segment processing. Fixed Block Chunking also provides a slightly better deduplication due to lower block size. In Veeam Backup & Replication version 12, the Fixed Block optimization can be used for both virtual and physical environment backup jobs generating full, synthetic full, and incremental (* .vbk and * .vib) backup files.

All Veeam Backup & Replication backup repositories created using Catalyst Stores with the **Enforce Fixed Block Chunking** option set will have **Align backup file data blocks** automatically enabled, as shown in Figure 5.

- **Storage Quotas** are used to set a limit on the current Catalyst Store size. These settings are particularly useful when an HPE StoreOnce appliance is shared by different groups or cost centers, and you need to avoid the situation where a group of clients consume more capacity than the desired maximum. The **Logical Storage Quota** is a set limit on the amount of data written by Veeam independently of the deduplication. The **Physical Storage Quota** sets a limit on the HPE StoreOnce capacity utilization and is affected by the deduplication.
- The **Session History** defines the duration for which the session logs should be retained in an HPE StoreOnce appliance.




Figure 5. “Align backup file data blocks” enabled for the Veeam Backup & Replication Catalyst backup repository



Note

Enforce Fixed Block Chunking is applicable only for Low Bandwidth (source-side deduplication) transfer policies. Fixed Block Chunking can only be enabled during Catalyst Store creation. It cannot be set at a later point of time. This is an intentional restriction to prevent the user to mix inside the same Catalyst Store data processed by different deduplication algorithms that could decrease the deduplication efficiency.

Configuring access permissions to an HPE StoreOnce Catalyst Store

After a Catalyst Store is created, Hewlett Packard Enterprise recommends configuring **Client Access** to allow a selected group of client identifiers to have access to the HPE StoreOnce Catalyst Stores. To create and configure access to a client, navigate to the **Permissions tab** of the required Catalyst Store and add a client as shown in Figure 6.

Figure 6. Creating a client and granting permissions to a Catalyst Store from HPE StoreOnce GUI

Take note of the **Catalyst Client Name** and **Password** while creating a client. These credentials are required when you create a backup repository using this Catalyst Store. This step is demonstrated in [Creating an immutable HPE StoreOnce Catalyst-based Veeam backup repository](#) section. For more details on setting Client Access see [Integrating backup applications with StoreOnce Catalyst clients overview](#).

How to connect to an HPE StoreOnce Catalyst Store via IP or Fibre Channel

There are two options for connecting Veeam servers to an HPE StoreOnce Catalyst Store.

1. **IP/Ethernet:** Catalyst over Ethernet (CoE) is the most commonly used option, and it is available with minimal configuration. This configuration is supported by Veeam gateways running on both physical and virtual servers. Every Veeam proxy server can also run the gateway service and directly access HPE StoreOnce via Catalyst protocol, thereby avoiding an additional hop in LAN. CoE is generally slightly faster than Catalyst over FC (CoFC).
2. **Fibre Channel (FC):** CoFC functions the same way as HPE StoreOnce CoE. However, some additional configuration is required to set up the backup and restore connections between the ports on the HPE StoreOnce appliance and the ports on the client servers. This is done using the **Fibre Channel** settings tab in the HPE StoreOnce GUI, shown in [Figure 7](#). This GUI page is only available if CoFC is enabled on the HPE StoreOnce appliance. CoFC is only supported by Veeam gateways running on Windows physical servers. The configuration is less flexible than Ethernet and it is generally recommended when there isn't enough Ethernet bandwidth, or when there is a requirement to use a dedicated and isolated connectivity.

Status	Port Location	Current Speed	Selected Speed	Beacon
Up (F Port)	Slot 5 Port 1	16G	Auto-negotiate	Off
Up (F Port)	Slot 5 Port 2	16G	Auto-negotiate	Off

Figure 7. Configuring FC access to an HPE StoreOnce Catalyst Store

Initial Fibre Channel configuration

The following are prerequisites for running CoFC backups:

- The Veeam gateway server and the HPE StoreOnce appliance are connected to the Fibre Channel network.
- The Windows server that will run the Veeam gateway service can communicate with the HPE StoreOnce appliance over a Fibre Channel network. To do so, the storage administrator must make sure that any network segregation, such as zoning, is set up to handle the required connectivity between server and appliance.

The FC settings of the HPE StoreOnce appliance are shown in Figure 8, which is located at **Settings → Catalyst Settings → Fibre Channel**.

Device Name	Status	FC Address	World Wide Node Name	Number of Devices per Login
Target Port 5	Online	6A0101	50:01:43:80:34:43:8c:f9	20

Figure 8. Catalyst “Fibre Channel” settings tab shows Fibre Channel settings and the number of target ports and FC devices

The following information should be obtained from this **Fibre Channel** settings tab:

- Locate the FC **Identifier** as shown in Figure 8. This is the Fibre Channel address of the HPE StoreOnce appliance that is configured in Veeam Backup & Replication. It is in the format “COFC-*<device-id>*,” which will be input into Veeam instead of an Ethernet address.
- Hewlett Packard Enterprise recommends using the default values in the target ports section.
- Locate the **World Wide Port Name** for each port on the HPE StoreOnce appliance. This is the information needed to zone the Veeam gateway server Fibre Channel ports with the HPE StoreOnce appliance Fibre Channel ports.
- HPE StoreOnce Catalyst Copy over Fibre Channel is a two-way protocol. Zone the source initiator WWN with the destination target WWN and zone the destination initiator WWN with the source target WWN. The source and destination must be able to communicate with each other over Fibre Channel.



- Locate the **Number of Devices per Login** for each port. This value determines the number of backup and restore connections that are allowed from each Veeam gateway server Fibre Channel port on each available path to the ports on the HPE StoreOnce appliance.
 - The number of possible FC connections between a Veeam server and an HPE StoreOnce appliance is calculated as:
 $\text{number of client ports zoned} * \text{number of HPE StoreOnce node ports zoned} * \text{number of devices per login}$.
 Table 1 illustrates how to count the maximum number of FC connections between the Veeam gateway server FC ports and the HPE StoreOnce appliance FC ports.
 - In most situations, this value should be increased to the maximum number of possible connections available in the Veeam Backup & Replication GUI. The only reason to limit the FC connection is when the HPE StoreOnce unit is shared among multiple protocols and multiple backup applications, and the administrators want to avoid a single protocol or backup application using too many resources.
 - If this value is changed, a device-file rescan on the Veeam gateway server is needed before the change is recognized as shown in [Figure 11](#).

Table 1. Example relationships between device per initiator port and number of data sessions

Veeam server	HPE StoreOnce appliance	Number of devices per login (user modifiable)	Resulting in the maximum number of FC connections between Veeam gateway server FC ports and HPE StoreOnce FC ports
1 port zoned to	4 ports	1	4
2 ports zoned to	4 ports	1	8
1 port zoned to	1 port	64	64
2 ports zoned to	2 ports	4	16
2 ports zoned to	4 ports	8	64

Deploying a Veeam gateway server

A Veeam gateway server hosts the data-mover component required for communicating with HPE StoreOnce. Usually, the Veeam gateway service and the Veeam proxy services run on the same physical or virtual server. The role of a gateway server can be assigned to a Windows machine and added as a component in the Veeam Backup server. For more information, see [Veeam Backup & Replication 12 - Gateway Server](#).

The following are prerequisites for a Veeam gateway server:

- Administrator privileges are required to access OS-specific device files associated with HPE StoreOnce CoFC devices.
- Fibre Channel users are advised to consult the [HPE StoreOnce Support Matrix](#) (SPOCK account required) to be sure that Veeam gateway server host bus adapters (HBAs), switches, Fibre Channel driver, and firmware versions are supported.

To configure a gateway server, in the Veeam Backup & Replication GUI, navigate to **Backup infrastructure** and click **Add Server** as shown in Figure 9.







Figure 9. Adding a Veeam Gateway server role to a Microsoft Windows server

Supply the DNS name or IP address of a Microsoft Windows server, as shown in Figure 10.





New Windows Server

Name
Specify DNS name or IP address of Microsoft Windows server.

Name
DNS name or IP address:
192.168.20.63

Description:
Created by HPELAB\grishma at 2/28/2023 5:38 AM.

< Previous Next > Finish Cancel

Figure 10. Adding a gateway server role to a Microsoft Windows server

Review and apply the settings to the Microsoft Windows server, and then click **Finish** to apply the settings.

Configuring the Veeam gateway server for CoFC

HPE StoreOnce CoFC presents a device type of “Processor.” These devices are shown as **Other devices** in Windows Device Manager. After zoning the devices or changing the number of devices per initiator port, right-click **Other devices** and click **Scan for hardware changes** to detect the new devices. This is demonstrated in Figure 11.






Figure 11. HPE StoreOnce CoFC presents a device type of Processor and is shown under “Other devices” in Windows Device Manager

Creating an immutable HPE StoreOnce Catalyst-based Veeam backup repository

A Veeam backup repository is simply defined as storage for Veeam backup files created by Veeam Backup & Replication. In the case of an HPE StoreOnce Catalyst Store, the Veeam backup proxy reads backup data from the source infrastructure and moves it to the Catalyst Store via a Veeam gateway server. To create a Catalyst-enabled Veeam backup repository, perform the following steps:

1. Click **Backup Infrastructure** in Veeam Backup & Replication.
2. Right-click **Backup Repositories** and select **Add backup repository** as shown in Figure 12.






Figure 12. Veeam GUI showing the creation of a new backup repository

3. Select **Deduplicating storage appliance** and select **HPE StoreOnce** as the **Deduplicating Storage Appliance**, shown in Figure 13.






Figure 13. Specifying “HPE StoreOnce” as the deduplication storage appliance

4. The HPE StoreOnce backup repository can now be configured. A name and an optional description can be given to the repository in the **Name** wizard section.
5. In the HPE StoreOnce wizard section, provide:

HPE StoreOnce server name –The HPE StoreOnce server name can be specified in two forms for CoE and CoFC:

 - For CoE, type IP or DNS name
 - For the DNS name, use a FQDN to avoid potential errors.
 - See Figure 14 as an example of CoE.



New Backup Repository ×

 **HPE StoreOnce**
Specify HPE StoreOnce storage name and credentials.


Name	Type in HPE StoreOnce server name:
HPE StoreOnce	192.168.16.23
Repository	☐ Use Fibre Channel (FC) connectivity Requires that gateway server is connected into SAN fabric.
Mount Server	Credentials:
Review	Puser (Puser, last edited: less than a day ago) Add...
Apply	Manage accounts
Summary	Gateway server:
	Automatic selection Choose...
	☐ Gateway server and StoreOnce are connected over WAN Enables network traffic compression and checksumming by Catalyst. Using this functionality may reduce backup performance over fast links.

< Previous Next > Finish Cancel

Figure 14. Specifying HPE StoreOnce server name for Catalyst over Ethernet (CoE)

- For CoFC, type the HPE StoreOnce FC identifier, which has the format “COFC-<device-id>”
See Figure 15 as an example of CoFC.





New Backup Repository

HPE StoreOnce
Specify HPE StoreOnce storage name and credentials.

Name
Type in HPE StoreOnce server name:
COFC-SGH230W269

☒ Use Fibre Channel (FC) connectivity
Requires that gateway server is connected into SAN fabric.

Credentials:
Puser (Puser, last edited: 17 days ago) [Add...](#) [Manage accounts](#)

Gateway server:
WIN-M5RVLC3PH3P.hpelab.local [Choose...](#)

☐ Gateway server and StoreOnce are connected over WAN
Enables network traffic compression and checksumming by Catalyst. Using this functionality may reduce backup performance over fast links.

< Previous Next > Finish Cancel

Figure 15. Specifying HPE StoreOnce server name for Catalyst over FC (CoFC)

As shown in Figure 14 and Figure 15, there are other important configuration fields:

Credentials—Credentials to access HPE StoreOnce are specified here. If **Client Access** has been enabled, those credentials will be used here. Otherwise, any user or admin credentials can be used. New credentials can be added by clicking **Add...** and entering the appropriate credentials that will be stored in the Veeam credentials database for future use. In this instance, you should use the same credentials of the client that has access to the Catalyst Store

Gateway server—This is a Windows server running the Veeam gateway server service. The gateway server reads and writes data to an HPE StoreOnce appliance. For more information on assigning a gateway server, see [Deploying a Veeam gateway server](#).

- For CoFC, the gateway server(s) that has FC connectivity with the HPE StoreOnce appliance must be selected. Click **Choose** as shown in Figure 15 and select the required server(s) listed under **Use the selected gateway servers only** option.
- For CoE, you can choose **Automatic selection** or specify the required servers under **Use the following gateway servers only** option as shown in [Figure 16](#).

Note

Starting from Veeam Backup & Replication Version 12, multiple gateway servers that are part of a gateway pool can write to the same backup repository for expediting the backup process. To enable the use of a gateway pool for a backup job, either select **Automatic selection** or select multiple gateway servers as shown in [Figure 16](#).

- Gateway server and StoreOnce are connected over WAN**—Select this checkbox when the communication between the Veeam gateway server and HPE StoreOnce appliance is across a WAN link. This optimization option enables Catalyst native compression and checksum algorithms. Compression further reduces the WAN bandwidth utilization and checksum makes the communication more resilient to communication errors that are more frequent on WAN links.

Note
Hewlett Packard Enterprise recommends configuring both proxy and gateway services to run on the same server whenever possible. This eliminates an extra hop in LAN for backup traffic and reduces backup proxy/gateway CPU utilization. Veeam Backup & Replication Version 12 should automatically choose the gateway on the same proxy server when available. For a manual configuration, to be sure that the gateway and proxy services are on the same server(s), you can specify the same server name(s) on **Use the following gateway servers only** as shown in Figure 16, and on the **Backup proxy:** in the **Backup Job** configuration wizard. See the [Creating a Veeam backup job](#) section for more information on the **Backup Job** configuration wizard.




Figure 16. Choosing required gateway servers

6. The **Repository** settings in the **New Backup Repository** wizard are illustrated in Figure 17.



New Backup Repository

Repository
Type in path to the folder where backup files should be stored, and set repository load control options.

Name
HPE StoreOnce

Repository
Mount Server
Review
Apply
Summary

Location
Catalyst store: CSStoreImmutability
Capacity: <Unknown> Free space: <Unknown> Populate
☒ Make recent backups immutable for: 14 days
Protects backups from modification or deletion by ransomware, malicious insiders and hackers. GFS backups are made immutable for the entire duration of their retention policy.

Load control
Running too many concurrent tasks against the repository may reduce overall performance, and cause I/O timeouts. Control storage device saturation with the following settings:
☐ Limit maximum concurrent tasks to: 4
☐ Limit read and write data rate to: 1 MB/s

Click Advanced to customize repository settings. Advanced...

< Previous Next > Finish Cancel

Figure 17. “New Backup Repository” wizard specifying Catalyst Store and load control settings

Aspects of the settings are:

- **Catalyst Store**—Select the Catalyst Store created in the [Creating an Immutable Catalyst store for Veeam backups](#) section to associate with this Veeam repository. To see the list of accessible Catalyst Stores, click the down-arrow at the right. Only the stores compatible with your provided credentials are shown.
- Select the **Make recent backups immutable for <> days** checkbox and specify the desired number of days for which the backup objects in the Catalyst Store need to be write-protected.

Note

The immutability duration set while creating or modifying HPE StoreOnce Catalyst repository should not be greater than the Maximum ISV controlled immutability duration specified during Catalyst Store creation. This value can be changed at any time. If necessary, you can check the Maximum ISV controlled immutability duration on the HPE StoreOnce appliance and set it to greater than or equal to 365000 days.

- **Load Control**— Load control should be kept disabled to maximize performance. Veeam Backup & Replication monitors the HPE StoreOnce resource utilization and never starts more sessions than the maximum number supported by the appliance model in use, so Load control is not necessary to be set. Load control can be enabled when it is necessary to limit the HPE StoreOnce resource consumption by a single Veeam Backup & Replication repository. This is required especially when there are multiple clients/applications sharing the same HPE StoreOnce appliance.
- **Advanced...**—Click the **Advanced** button to configure **Storage Compatibility Settings** as illustrated in Figure 18.






Figure 18. "Storage Compatibility Settings" in the "New Backup Repository" wizard

Note

Do not change the default setting as shown in Figure 18. Advanced settings should not be adjusted for HPE StoreOnce-based backup repositories.

Storage compatibility settings can be configured as follows:

- **Align backup file data blocks**—This should be enabled along with enforcing Fixed Block (FB) Chunking at the Catalyst Store level for better HPE StoreOnce deduplication. This is the recommended setting for all Veeam workloads supporting Fixed Block Chunking. Backup file processing with these settings is found to improve per-stream backup performance by 4 to 5 times. It also improves synthetic full backup performance and reduces HPE StoreOnce storage utilization. When Fixed Block Chunking is not enabled at the Catalyst Store level, uncheck **Align backup file data blocks** to minimize storage utilization. For more details on Fixed Block Chunking, see the [Catalyst Store Advanced Settings](#) section.
- **Decompress backup data blocks before storing** —This should be enabled for better Catalyst deduplication. Hewlett Packard Enterprise recommends keeping compression set to **Optimal** on the backup job and **Decompress backup data blocks before storing** on the backup repository for better performance and capacity utilization.
- **Use per-machine backup files** —Starting from version 12 of Veeam Backup & Replication, this option is enabled by default for HPE StoreOnce backup repositories. The setting enables the creation of one backup file per VM, which gets mapped to a Catalyst I/O stream thereby improving overall backup performance. This also results in a true per-VM backup chain.



7. The **Mount Server** settings in the **New Backup Repository** wizard are illustrated in Figure 19.

New Backup Repository

Mount Server
Specify a server to mount backups to when performing advanced restores (file, application item and instant VM recoveries). Instant recoveries require a write cache folder to store changed disk blocks in.

Name
HPE StoreOnce

Repository

Mount Server

Review

Apply

Summary

Mount server:
veeamsvr1.hpelab.local (Created by HPELAB\grishma at 3/4/2023 7:45 AM.) **Add New...**

Instant recovery write cache folder:
C:\ProgramData\Veeam\Backup\IRCache\ **Browse...**

Ensure that the selected volume has sufficient free disk space to store changed disk blocks of instantly recovered machines. We recommend placing the write cache folder on an SSD drive.

☒ **Enable vPower NFS service on the mount server (recommended)** **Ports...**
Unlocks instant recovery of any backup (physical, virtual or cloud) to a VMware vSphere VM.
vPower NFS service is not used for instant recovery to a Microsoft Hyper-V VM.

< Previous **Next >** **Finish** **Cancel**

Figure 19. Configuring the mount server for the Veeam backup repository

Mount Server settings can be configured as follows:

- **Mount server**—A mount server is used to mount backup files for Instant VM Recovery, SureBackup, and On-Demand Sandbox™. For granular restores, the backup server will automatically play this role. In the case of HPE StoreOnce, in most situations, it is a good practice to select the same server for the mount server and the gateway server to avoid an extra hop in LAN.
- **Instant recovery write cache folder**—This is used as a temporary location to store the write cache for machines that are started from backups during recovery verification or restore operations. For more details, see the [Veeam Backup & Replication 12 User Guide](#).
- **vPower NFS**—When a Veeam Instant VM Recovery is performed, backups are mounted to the hypervisor and powered on in a read-only state. All new data will be written to a temporary cache stored at the path specified in Figure 19.

8. Review all settings and then click **Finish**.

Creating multiple Veeam backup repositories inside the same HPE StoreOnce Catalyst Store

There are situations where it is useful to create multiple Veeam backup repositories inside the same Catalyst Store. The common reasons for doing this are:

- **Better deduplication**—Each Catalyst Store is an independent deduplication domain. To enable cross-deduplication among multiple backup repositories, it is possible to create them inside the same Catalyst Store. This is useful when we backup similar data to different backup repositories.
- **Catalyst Copy granularity**—As described in the [Veeam-managed HPE StoreOnce Catalyst Copy job](#) section, the Veeam Catalyst Copy job copies the contents of an entire backup repository to other HPE StoreOnce appliances. When multiple jobs write to the same backup repository, the Catalyst Copy job will replicate the backup data of all backup jobs. There are situations where it is necessary to tailor the replication parameters to the systems protected by specific jobs. This configuration requires multiple backup repositories—potentially one per job—and it can be useful to have them in the same Catalyst Store to get better deduplication.
- **Migration**—Veeam provides an easy methodology for migrating entire backup repositories to new storage platforms. (See the [Migrating Veeam backup repositories to/from an HPE StoreOnce Catalyst Store](#) section for details.) If a storage platform becomes full, and you want to migrate a subset to new storage, then a solution design based on multiple backup repositories offers more flexibility than a solution based on a single large backup repository.



- Manual workload balancing—Starting with Veeam Backup & Replication Version 12, manual workload balancing is generally not necessary. Veeam Backup & Replication version 12 has an effective dynamic load balancing mechanism that can use multiple gateways to concurrently access the same backup repository. Veeam Backup & Replication will try to use the gateway service running on the same server running the proxy service. This optimization is useful to better distribute the workload for VMs inside the same job, assigning them to different proxies and gateways, and at the same time, to avoid an extra hop in LAN that would occur when the proxy and gateway are on different servers.

If a manual workload balancing is required, it is possible to configure a static balancing, defining which proxy will run a job and which gateway will manage the backup repository. In this situation, it is important to assign the proxy and gateway services to the same server. This configuration requires multiple backup repositories, potentially one per job, and it can be useful to have them in the same Catalyst Store to get better deduplication.

Important

Veeam recommends not using the Veeam GUI for creating multiple backup repositories (BRs) on the same Catalyst Store. The configuration made using the GUI can create instability and should be avoided. Veeam recommends using the following supported PowerShell commands for creating multiple BRs instead of using the GUI. After its created, the backup repository can be managed by the GUI, except for the **Repository** tab. For more details, see [veeam.com/kb2987](https://www.veeam.com/kb2987).

After it is created by the CLI process, a backup repository can be conveniently managed by the GUI as with any other backup repository. The following examples demonstrate how to properly set up Veeam Backup Repositories via the CLI.

To open the Veeam PowerShell screen, click the icon on the top-left Veeam GUI corner, select **Console**, and then click **PowerShell**, as shown in Figure 20.




Figure 20. How to open Veeam PowerShell command window

Example 1—Create the additional “BR01” repository in the “MultiBR” Catalyst Store:

```
Add-VBRBackupRepository -Name BR01 -Folder storeonce://so-vsa-01.hpelab.local:ImmutableFBStore@/repo01
-Type HPStoreOnceIntegration -StoreOnceServerName so-vsa-01.hpelab.local -UserName veeam -Password xxxxxx
```



Example 2—Remove the limits for “maximum concurrent tasks” and for “read and write data rate.” This change is only possible via the CLI.

```
$repository = Get-VBRBackupRepository -Name "BR01"

set-VBRBackupRepository -Repository $repository -DataRateLimit 12
-LimitDataRate:$false -MaxConcurrentJobs 12 -LimitConcurrentJobs:$false
```

Note

The procedure described for creating multiple backup repositories is applicable for HPE StoreOnce Cloud Bank Catalyst Stores as well.

Creating a primary backup job for VMs

A Veeam backup job acts as a policy for protecting one or more VMs. A backup job holds specific settings for compression, deduplication, block transmission size, application quiescence, and execution schedules. A backup job also defines the backup mode, such as forward incremental backups with periodic full backups. There are many types of Veeam backup jobs that support HPE StoreOnce Catalyst. See [Veeam Backup & Replication 12 - HPE StoreOnce](#) for backup deployment options.

This section provides the recommended settings when creating a backup job writing to an HPE StoreOnce Catalyst-based backup repository. To add VMs to a backup job, perform these steps:

1. In the Veeam Backup & Replication console, click **HOME** and then **Jobs** to see the **Backup Job** selection in the ribbon, as demonstrated in Figure 21.




Figure 21. Navigating to a Veeam Backup Job

2. Using the ribbon menu at the top of the user interface, click **Backup job**, then select **Virtual machine** as the Backup job type
3. In the **Name** wizard section, give the job a meaningful **Name** and optional **Description**, and then click **Next**.



4. On the **Virtual Machines** wizard section, click **Add** and select the VMs to add to this backup job, as demonstrated in Figure 22.




Figure 22. Adding VMs to a backup job in the Veeam GUI

Note

If a hierarchy object such as “Resource Pool” or “Cluster” is chosen, Veeam will back up all VMs under this object. Also, new VMs added under the object will be backed up on subsequent backup job runs or removed from the job if they are deleted from the object.

Tip

If the infrastructure is designed for storage snapshot integration (best practice in all Veeam Backup infrastructures), it is generally useful to use the “Volume” view and select entire volumes.

5. The **Storage** wizard section displays the configuration elements shown in Figure 23.



New Backup Job



Storage

Specify processing proxy server to be used for source data retrieval, backup repository to store the backup files produced by this job and customize advanced job settings if required.

Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Backup proxy:

Automatic selection

Choose...

Backup repository:

gri-corgi-bkprepo1 (Created by HPELAB\grishma at 3/3/2023 4:36 AM.)

146 TB free of 147 TB

Map backup

Retention policy: 7 days

☐ Keep certain full backups longer for archival purposes
GFS retention policy is not configured

☐ Configure secondary destinations for this job
Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.

Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings.

Advanced...

< Previous

Next >

Finish

Cancel

Figure 23. Specifying proxy, repository, and retention parameters in the “New Backup Job” wizard

6. After adding VMs to the backup job, the storage settings shown in Figure 23 can now be specified. Storage settings can be summarized as follows:

- **Backup proxy**—Veeam’s intelligent load balancing between multiple backup proxies can be leveraged by leaving this as **Automatic**. Whenever possible, it is recommended that the backup proxy server be the same as the gateway server selected in the **Backup Repository** configuration wizard. This will reduce the number of hops in the data path and increases performance.

Note

Veeam Backup & Replication version 12 can execute the gateway service on multiple servers concurrently, even for writing to the same backup repository on a HPE StoreOnce Catalyst Store. This makes it easier to create configurations that automatically offer load balancing, avoid bottlenecks, and reduce the occurrence of unnecessary additional hop in LAN for the backup data.

- **Backup repository**—Select the HPE StoreOnce Catalyst-based backup repository created in previous steps.
- **Retention policy**—The retention policy can be specified in **restore points** or **days**. The frequency of the Veeam backup job and immutability retention period set at the HPE StoreOnce backup repository are important aspects to consider while defining the retention policy.
- **Keep certain full backups for longer for archival purposes**—The long-term or Grandfather-Father-Son (GFS) retention policy allows you to store selected backup files for extended periods of time, including weeks, months, and years. Veeam Backup & Replication does not need to create any new backup files—it uses backup files already created and marked them with specific GFS flags. For Catalyst-based backup repositories, GFS retention policies are generally more efficient than using a Backup Copy Job. For more information, see [Veeam Backup & Replication 12 - Long-Term Retention Policy \(GFS\)](#).



Note

HPE StoreOnce deduplication greatly reduces the storage requirement for additional restore points and longer daily retention. For example, in many production backup environments, when the retention is doubled from two to four weeks, the storage utilization on HPE StoreOnce does not double but grows only by ~19%. If retention is extended from two to eight weeks, the storage utilization on HPE StoreOnce grows only by ~57%. These are average values and in most environments the differences are minimal, but the final deduplication effect can vary, and it is highly influenced by the protected data type.

For better deduplication, whenever possible, Hewlett Packard Enterprise recommends writing backup data from similar VMs to the same Catalyst Store. This is explained in more detail in the [Creating multiple Veeam backup repositories inside the same HPE StoreOnce Catalyst Store](#) section. This can be achieved in several ways:

- Group VMs with similar data in the same job.
- Use the same backup repository for multiple jobs containing VMs with similar data.
- Create multiple backup repositories writing to the same Catalyst Store and use those repositories for multiple jobs containing VMs with similar data.

7. Clicking the **Advanced** button in [Figure 23](#) opens the backup job **Advanced Settings**, as shown in Figure 24

The screenshot shows the 'Advanced Settings' dialog box for a backup job. The 'Backup mode' section has two radio buttons: 'Reverse incremental (slower)' and 'Incremental (recommended)'. The 'Incremental (recommended)' option is selected. Below it, there is a checkbox 'Create synthetic full backups periodically on:' which is checked, with 'Saturday' selected and a 'Configure...' link. The 'Active full backup' section has a checkbox 'Create active full backups periodically on:' which is unchecked, with 'Saturday' selected and a 'Configure...' link. A 'Schedule Settings' sub-dialog box is open, showing 'Schedule:' with 'Monthly on:' and 'Weekly:' options. 'Weekly:' is selected, and 'On these days' is chosen from the dropdown. 'Saturday' is listed under 'On these days'. The 'OK' button is highlighted. The main dialog has tabs for 'Backup', 'Maintenance', 'Storage', 'Notifications', 'vSphere', 'Integration', and 'Scripts'. The 'Backup' tab is active. At the bottom of the main dialog are 'Save As Default', 'OK', and 'Cancel' buttons.

Figure 24. Specifying backup mode and synthetic full backups in the “Advanced Settings” of the “New Backup Job” wizard

8. The settings in the **Backup** tab shown in Figure 24 should be configured as follows:

- **Backup mode**—Incremental
- **Synthetic Full backups**—The synthetic full backup creation process with HPE StoreOnce is highly optimized and greatly reduces the full backup workload. The Veeam Backup & Replication incremental backup policy requires to “Create synthetic full backup periodically”, or to “Create active full backup periodically”. This full backup keeps the backup chain short and makes the backup and restore process more efficient. The synthetic full backup process offers a great advantage over the traditional full backup as it reduces the workload on the production infrastructure by 90%. Veeam leverages the Catalyst API to offload this process from the Veeam Backup & Replication servers to the HPE StoreOnce appliance. The HPE StoreOnce appliance performs this task very efficiently as it does not move the actual data, but only reorganizes pointers to the backup data (the latest backup chain) that is already persisted in the system.

This process of synthetic full backup creation is described as follows:

- a. Veeam backup job creates an incremental backup. A full backup is not created on the day of synthetic full backup creation, and this reduces about 90% of the workload on production storage and servers.
- b. Veeam sends the “instructions” to the HPE StoreOnce appliance for merging the previous full with all the subsequent incremental backups to finally generate the new synthetic full backup.
- c. HPE StoreOnce executes this task at high speed as it does not need to move the actual data, but only the pointers to the stored deduplicated data segments. To enable synthetic full backups, select the checkbox **Create synthetic full backup periodically**, and choose when this should occur. A synthetic full backup generally completes in 50% or less of the time required for a traditional active full backup.
- d. Starting with Veeam Backup & Replication Version 12 and HPE StoreOnce 4.3.2, there are additional performance improvements in the synthetic full backup process.
 - i. The synthetic full transformation is performed immediately upon completion of incremental backup at the individual VM level rather than waiting for the completion of incremental backup of all the VMs associated with the backup job. This reduces the overall backup window and avoids the workload peak previously had at the end of a job execution when the synthetic full operation was triggered for all the VMs at the same time.
 - ii. Fixed Block Processing for Veeam, along with the improvement in metadata reconstruction in the HPE StoreOnce firmware 4.3.2, significantly increases the synthetic full backup performance.

Note

A synthetic full backup generally completes in 50% or less of the time required for a traditional active full backup, and the workload for the production environment is about 90% lower. Because of these advantages, Hewlett Packard Enterprise recommends using the **Create synthetic full backup periodically**, rather than the **Create active full backup periodically** option.



9. The settings in the **Maintenance** tab shown in Figure 25 should be configured as follows:




Figure 25. "Maintenance" settings are not required for Catalyst-based repositories

- **Storage-level corruption guard** settings should be left unchecked. HPE StoreOnce has internal data integrity checking and self-healing mechanisms. Enabling this option could lead to severe performance degradation during backup process.
- **Full backup file maintenance** settings must be left unchecked because the backup policy already includes periodic full backup.



10. For better performance and deduplication, Hewlett Packard Enterprise recommends the following **Storage** advanced settings, as illustrated in Figure 26.

Advanced Settings

Backup Maintenance **Storage** Notifications vSphere Integration Scripts

Data reduction

☐ Enable inline data deduplication (recommended)

☒ Exclude swap file blocks (recommended)

☒ Exclude deleted file blocks (recommended)

Compression level:

Optimal (recommended)

Provides for the best compression to performance ratio, lowest backup proxy CPU usage and fastest restore.

Storage optimization:

4MB

Required for processing machines with disks larger than 100TB. Reduces dedupe ratio and increases the size of incremental backups.

Encryption

☐ Enable backup file encryption

Password:

Manage passwords

Save As Default OK Cancel

Figure 26. Specifying deduplication, compression, optimization, and encryption options

Data deduplication—Disabled

Compression level—“Optimal”

Storage Optimization—“4MB”

Encryption—Disabled

Note

If data-at-rest encryption is required, use HPE StoreOnce embedded encryption rather than Veeam encryption because Veeam encryption will highly reduce HPE StoreOnce data deduplication. For more information, see the [Creating a Catalyst Store for Veeam backups](#) section.

11. The **Integration** tab in backup job advanced settings (Figure 27) can be used to configure backups from HPE Primary Storage snapshots rather than from the VMware host as a best practice. To enable backups from HPE Primary Storage snapshots, additional configuration is required. More information can be found in the [HPE Reference Configuration for Veeam Availability Suite with HPE Primera Storage](#) and the [HPE Reference Configuration for Veeam Availability Suite with HPE Nimble Storage](#).

The use of storage snapshots within a backup process considerably reduces the workload for the production infrastructure and reduces the impact that the backup has on production response time. It also greatly reduces the risk of a “stun” effect on production services that can occur when long-lasting VM snapshots are deleted, and a lot of data has to be consolidated to the main vmdk file.






Figure 27. Configuring backups from storage snapshots

Note

Veeam attempts to concurrently process all the VMs in a storage snapshot. If necessary, the number can be limited. See Figure 27 for **Limit processed VM count per storage snapshot to**. In general, it is better to process all the VMs concurrently to avoid taking multiple HW snapshots. With HW snapshot integration, the overhead of processing multiple VMs concurrently is limited because each VM snapshot has a short life, and consequently, a lower impact on production infrastructure.

12. On the **Guest Processing** tab of the Veeam Backup Job wizard, it is possible to enable transaction consistency for VM backups and indexing of VM guest file systems for quick searching through the optional Veeam Backup Enterprise Manager. To enable these capabilities, simply check the appropriate box and supply administrator credentials. Additional details on the configuration of these features can be found at [Veeam Backup & Replication 12 - Step 10. Specify Guest Processing Settings](#).
13. On the **Schedule** tab, it is possible to control the backup job scheduling. Usually, backup jobs are run daily during the night. If more frequent recovery points are required, a more frequent synthetic full backup schedule will need to be used to keep the backchain at a reasonable length.
14. View **Summary** and create a new schedule according to the recommended settings.

Creating a Veeam backup job for Windows and Linux computers

Veeam Agents can also be utilized to backup data from Windows and Linux® computers to an immutable HPE StoreOnce Catalyst Store target, which can yield the same level of protection and deduplication of virtual machine data.

Backup data from a Veeam Agent can be transferred to HPE StoreOnce in two ways. The first is the server-based backup, demonstrated in Figure 28, which supports both Windows and Linux servers, and requires a Veeam gateway service to write to HPE StoreOnce.






Figure 28. Veeam Agent server-based backup

The second is the direct-client backup, illustrated in Figure 29, which supports only Windows servers. This backup is generally faster because it removes the extra hop in LAN. This is the suggested configuration for large clients where performance is important.




Figure 29. Veeam Agent direct-client backup

For configuring the Direct backup there are a few simple additional configuration steps:

1. Add each Windows Client to the Veeam-managed server list, as shown in Figure 30. This allows Veeam to install the gateway service for the next step. There is no need to make it a proxy.




Figure 30. Adding a Windows client as a Veeam-managed server



2. Starting from Veeam Backup & Replication Version 12, multiple gateway servers can write to the same HPE StoreOnce Catalyst backup repository. To enable direct backup of multiple Windows clients to the same Catalyst repository, select all the required Windows clients as the gateway servers, as shown in Figure 31. For a full description, see [Creating an HPE StoreOnce Catalyst-based Veeam backup repository](#).




Figure 31. Adding a new backup repository by using the Windows clients as the Veeam gateway servers

With the above configuration, Veeam Backup & Replication Version 12 automatically chooses the windows servers associated with the Veeam Agent Backup job as gateway servers to write to the same Catalyst backup repository in a local or remote HPE StoreOnce appliance with all the benefits of direct communication, high level of source-side deduplication, and network bandwidth reduction.

3. Finally, create a Veeam Agent Backup Job for the desired Windows clients and assign the associated Backup Repository (created in Step 2) as the backup target. For more details on how to configure a Veeam Agent Backup Jobs, see [Veeam Backup & Replication 12 - Creating Veeam Agent Backup Jobs](#).

Make sure these aspects are followed while creating new Agent Backup Job:

- The Immutable HPE StoreOnce backup repository is configured as described in the [Creating an HPE StoreOnce Catalyst-based Veeam backup repository](#) section and is selected as the target in the **Storage** wizard screen.
- Create an immutable HPE StoreOnce backup repository as described in the [Creating an HPE StoreOnce Catalyst-based Veeam backup repository](#) section. Select this immutable HPE StoreOnce Catalyst repository as the backup repository in the **Storage** wizard screen while creating a new Agent Backup Job.
- In the **Advanced** settings of the **Storage** wizard screen, the following settings are recommended:
 - a. In the **Backup** tab, enable **Create synthetic full backups periodically** and define how often they should occur. If you have multiple jobs, distribute the synthetic full on different weekdays to avoid workload peaks on the HPE StoreOnce appliance. Do not enable **Create active full backups periodically** because this is not needed for HPE StoreOnce and creates an unnecessary additional workload.
 - b. In the **Maintenance** tab, do not enable **Storage-level corruption guard** as HPE StoreOnce has internal data integrity checking and a self-healing mechanism. Also, do not enable **Full backup file maintenance** because the backup policy already includes periodic full backups.
 - c. Do not enable any of the settings in the **Storage** tab: for **Compression level**, select **Optimal**, for **Storage Optimization**, select **4MB** because it is faster and does not use more HPE StoreOnce capacity, and for **Encryption**, do not select **Enable backup file encryption**. If you want encryption for data-at-rest, use HPE StoreOnce encryption as described in the [Creating a Catalyst Store for Veeam backups](#) section.
 - d. In the **Integration** tab, optionally select **Enable backup from storage snapshots** for expedited backup process and to minimize the guest OS quiesce window.

Validation of backup immutability

All the restore points written by a backup job to a Catalyst-based backup repository with immutability setting (see Figure 17 for details) are immutable for the defined period. The immutability period, once set by Veeam, is then enforced by HPE StoreOnce. An immutable backup cannot be deleted or modified either from the Veeam Backup & Replication console or the HPE StoreOnce GUI. The immutability period expiration is visible in the Veeam Backup & Replication Backup Properties windows as shown in Figure 32.

Backup Properties VMBackupJob1 (ImmutableRepo)

Objects:

Name	Original Size
Prod-ubuntu-vm2	100 GB

Total size: 100 GB

Restore points:

Date	Type	Status
Restore points: 0		

Files:

Name	Data Size	Backup Size	Deduplication	Compression	Date	Immutable Until
Prod-ubuntu-vm2.vm-8463D2022-0...	18.2 GB	18.2 GB	1.0 x	1.0 x	5/26/2022 8:28:50 AM	6/9/2022
Prod-ubuntu-vm2.vm-8463D2022-0...	100 GB	18.2 GB	5.4 x	1.0 x	5/26/2022 5:38:58 AM	6/9/2022
Prod-ubuntu-vm2.vm-8463D2022-0...	100 GB	18.2 GB	5.4 x	1.0 x	5/26/2022 5:34:55 AM	6/9/2022
Prod-ubuntu-vm2.vm-8463D2022-0...	100 MB	101 MB	1.0 x	1.4 x	5/26/2022 5:32:56 AM	6/9/2022
Prod-ubuntu-vm2.vm-8463D2022-0...	100 MB	101 MB	1.0 x	1.4 x	5/26/2022 5:28:03 AM	6/9/2022
Prod-ubuntu-vm2.vm-8463D2022-0...	128 MB	129 MB	1.0 x	1.3 x	5/26/2022 5:17:50 AM	6/9/2022
Prod-ubuntu-vm2.vm-8463D2022-0...	100 GB	18.2 GB	5.4 x	1.0 x	5/26/2022 3:55:53 AM	6/9/2022

Backup size: 73.2 GB

Copy path

Close

Figure 32. "Immutable Until" field populated against restore points

HPE StoreOnce immutability protection on Veeam Backup & Replication-initiated delete operations

Any accidental or deliberate attempt to delete the backup objects from Veeam Backup & Replication GUI will result in an error as shown in Figure 33.



Name:	Backup Deletion Job	Status:	Warning
Action type:	Backup Deletion	Start time:	6/10/2022 7:03:11 AM
Initiated by:	HPELAB\grishma	End time:	6/10/2022 7:03:27 AM
Log			
Message	Dur...		
✓ Starting backup deletion job			
✓ Preparing objects for deletion			
✓ Building tasks list			
✓ Processing backup 1 out of 1 (100% done)	0:00...		
⚠ [VMBkpJob] Backup deleted with warning	0:00...		
⚠ Unable to delete 13 immutable backup files			
⚠ Backup files can be deleted after 6/23/2022 10:35 PM			
✓ ImmutableLocalFBRepo: 0 deleted, 0 skipped, 1 warned, 0 failed			
⚠ Job finished with warning at 6/10/2022 7:03:27 AM			

Figure 33. Deletion of restore points blocked until the expiration of immutability duration

There are multiple ways for deleting restore points from Veeam Backup & Replication console, and no matter which one is selected, the HPE StoreOnce immutability enforcement will make it impossible for Veeam to delete any restore point with unexpired immutability.

HPE StoreOnce immutability protection on HPE StoreOnce-initiated delete operations

The immutability protection also works for any disruptive operation initiated by an HPE StoreOnce administrator using the HPE StoreOnce GUI or CLI. When HPE StoreOnce is in Dual Authorization mode—a prerequisite for using immutability with Veeam Backup & Replication Version 12—there are multiple restricted operations. For the complete list, see [HPE StoreOnce Dual Authorization requests](#). Here are two common situations where the immutability helps against a potential cyber-attack when a malicious actor uses an HPE StoreOnce administrator account to delete restore points

1. **The attacker tries to delete a specific backup file:** The immutable lock set by the Catalyst APIs on Veeam backup files (.vbk and .vib) protects the primary backups from any modification. A malicious attempt by an attacker to delete the backup objects in the HPE StoreOnce Catalyst Store using the hacked HPE StoreOnce Admin user credentials is fended off as demonstrated in Figure 34.



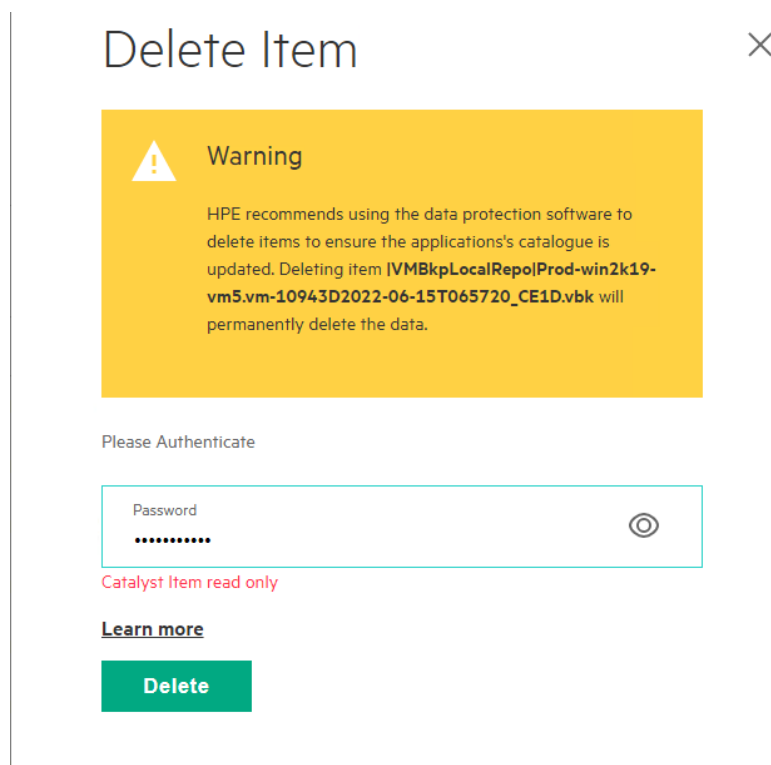


Figure 34. Deletion of backup object blocked by the immutable lock set by Catalyst APIs

2. **The attacker tries to delete an entire HPE StoreOnce Catalyst Store:** HPE StoreOnce provides the ability to delete an entire store and its content, including the immutable backup it may contain. However, the operation is blocked by Dual Authorization and it is not immediately executed. The deletion remains in a pending state until a security officer approves or denies it. For more details on Dual Authorization, see [Enabling Dual Authorization for HPE StoreOnce](#). Figure 35 shows an example of Catalyst Store modify and delete requests initiated by an HPE StoreOnce administrator that are queued for approval by the security officer.






Figure 35. Deletion of Catalyst Store queued for approval by the security officer

Copying backups to an off-site HPE StoreOnce and Cloud Bank Storage

Maintaining multiple copies of backups is critical to an organization's data availability strategy. The use of Backup Copy Jobs plays a key part in complying with the 3-2-1-1 rule: three copies of data on two different media types, with at least one off-site copy, and one immutable copy. Veeam Backup & Replication with HPE StoreOnce enables this by distributing/replicating restore points on geographically separated HPE StoreOnce appliances or to object storage located on-premises or in the cloud.

The following options are available for replicating/copying backups. For each solution design in the below list, a detailed description is also provided in the following chapters.

1. Veeam-managed HPE StoreOnce Catalyst Copy job (CCJ) to an off-site HPE StoreOnce Catalyst Store based repository—This is a storage-based, Veeam Backup & Replication controlled data replication from an HPE StoreOnce Catalyst Store to another Catalyst Store. The target Catalyst Store can be on the same, or more commonly, on a different HPE StoreOnce in LAN or WAN. Veeam Backup & Replication controls which backup file has to be replicated, the replica destination, and when the replication process occurs. The replication process is fully executed by a direct HPE StoreOnce-to-HPE StoreOnce communication. It does not require any additional server as a data mover and offers maximum efficiency, speed, and bandwidth reduction.
2. Veeam-managed Catalyst Copy Job (CCJ) to HPE StoreOnce Cloud Bank Storage—Using an HPE StoreOnce Cloud Bank Storage as a backup replication target, it is possible to create an off-site backup copy in the cloud. The replication process requires a fraction of the normally required capacity (and cost) in the object storage backing the HPE StoreOnce CBS Catalyst Store as only fully deduplicated and compressed backup data is transferred.
3. Veeam Backup Copy Job with Catalyst writes over LAN/WAN (BCJ) to an HPE StoreOnce repository—This is a server-based data replication to a local/remote HPE Catalyst-based repository. It requires one single server as a data mover and supports any source repository, such as a local file system. It requires an HPE StoreOnce appliance at the local/remote destination. This design is often used for small branch offices that do not have a local HPE StoreOnce appliance and want to replicate data to a central HPE StoreOnce destination. The process offers excellent bandwidth reduction.
4. Traditional Veeam Backup Copy Job configuration—This is a server-based data replication from/to any type of remote repository. It requires two servers as data movers. This design does not leverage any HPE StoreOnce bandwidth reduction, and it is not often used when the destination is an HPE StoreOnce appliance because the designs described above offer more advantages.

Creating a Veeam-managed Catalyst Copy Job to an HPE StoreOnce Catalyst Store-based repository

HPE StoreOnce Catalyst Copy operations can be conducted directly through Veeam. Veeam sends Catalyst Copy commands to an HPE StoreOnce appliance, and then HPE StoreOnce replicates the backup files from a source HPE StoreOnce appliance to an off-site HPE StoreOnce appliance. This scenario is illustrated in Figure 36.






Figure 36. CCJ using source and off-site HPE StoreOnce repositories

This configuration offers multiple advantages:

- Source and target backup repositories are shown on the Veeam GUI and restore points are visible and directly accessible for any restore operation similarly to any other type of repository. Users can easily configure and monitor the process from the Veeam GUI.
- Data is never rehydrated because the design does not require a Veeam server to read the data that has to be transferred. This removes the need to deploy an additional server and greatly reduces the workload on the source HPE StoreOnce appliance because it does not need to rehydrate any data.
- Parallel processing—HPE StoreOnce native replication has a wide queue length optimized for the WAN latency. This optimization increases the replication throughput, even when there is a single large file to be replicated.
- Best bandwidth reduction—data is replicated with the best possible deduplication and compression level. After the first backup replication, the average data reduction for each new data transfer is about 17:1 for incremental and 72:1 for a full backup. These are average values and in most environments the differences are minimal, but the deduplication effect can vary, and it is highly influenced by the protected data type.
- It does not require any additional data manipulation, such as rolling-full merge operations or periodic defragmentation. Data integrity checks are possible, but they are not required as HPE StoreOnce replication already has built-in internal consistency and corruption controls.

Veeam-managed CCJ also supports multi-site replication with parallel replication and cascade replication. This means that any source repository, including source repositories that are a target in a different CCJ, can be propagated to a different target repository. Figure 37 shows the multiple replication models that can be used for Veeam-managed HPE StoreOnce Catalyst Copy jobs.




Figure 37. Replication model involving multiple parallel and cascading Veeam-managed HPE StoreOnce Catalyst Copy jobs



The following best practices are recommended when configuring a Veeam-managed HPE StoreOnce Catalyst Copy job:

- Either Catalyst Copy over Ethernet (CCoE) or Catalyst Copy over Fibre Channel (CCoFC) can be used.
- If encryption for data-in-transit is required, if possible, configure your WAN network devices for making the point-to-point encryption. If this encryption facility is not available, then it is possible to configure data-in-transit encryption on HPE StoreOnce (see the [Creating a Catalyst store for Veeam backups](#) section), but this increases the workload and could reduce the replication performance.

To initiate a Veeam-managed HPE StoreOnce Catalyst Copy job, select **Backup Copy → Storage copy**




Figure 38. Initiating a Veeam-managed Catalyst Copy Job

After naming the job, in the **Storage Mapping** section, click **Add...** Designate the **Source backup repository** and the **Target backup repository** as shown in Figure 39.






Figure 39. Selecting the source and target HPE StoreOnce Catalyst Stores for a Veeam-managed HPE StoreOnce Catalyst Copy Job

Note

The replication protocol is defined by the target backup repository. If CCoE is desired, make sure to define the destination repository as CoE and not CoFC. It is supported to have the source repository accessed via CoFC and the destination via CoE. If the destination needs to be accessed via CoFC for performing restore operations, but the replication must use CCoE, then it is necessary to implement the following workaround:

1. The production Veeam instance defines the target repository via CoE.
2. Install a new Veeam server to be used for restoring the destination replica via CoFC. The new Veeam server can be a VM.
3. On the new Veeam server, create a backup repository that points to the replicated copy via CoFC, which requires a server with FC connectivity. Both the production server and the new Veeam server can use the same proxy/gateway, so you may not need to install a new physical server just for restoring via CoFC.
4. Rescan the backup repository to import the restore point list to the GUI.
5. The new Veeam server can be used for making a restore via CoFC from the Catalyst Copy replication destination.

For the advanced **Maintenance** settings shown in [Figure 40](#), disable the **Perform backup files health check** checkbox. When it is enabled, it increases the workload on the target HPE StoreOnce system, and the replication process takes longer to complete. HPE StoreOnce contains a sophisticated mechanism to ensure that the replicated file is the exact copy of the source file.

By default, the secondary copies have a grace expiration period of seven days when compared to the primary copies. To modify this behavior, select the **Keep secondary copies for additional days** checkbox and specify the desired retention time.






Figure 40. Veeam-managed Catalyst Copy job advanced "Maintenance" settings

Like other Backup Copy Jobs, this job can be run on a schedule. This is configured in the **Schedule** section of the Backup Copy Job, shown in Figure 41.



New Backup Copy Job

Schedule
Specify the backup copy job schedule.

Job
Storage Mapping
Schedule
Summary

This job can transfer data:
☒ Any time (continuously)
☐ During the following time periods only:

All None

12 2 4 6 8 10 12 2 4 6 8 10 12
 AM PM

Sunday
Monday
Tuesday
Wednesday
Thursday
Friday
Saturday

☒ Enabled ☐ Disabled

< Previous Apply Finish Cancel

Figure 41. Veeam-managed Catalyst Copy Jobs can be run on a schedule

File replication behavior from a Catalyst Copy Job (CCJ) can be described as follows:

- A CCJ makes a mirror of the source backup repository. If a file is deleted from the source, then the delete operation is replicated to the destination based on the secondary copy retention value set in [Figure 40](#). However, a completely unique GFS retention schedule is not configurable.
- A CCJ only replicates the backup files generated by a Backup Job or Backup Copy Job for VMware, Hyper-V, Nutanix AHV, RHV, and Managed/Unmanaged Agents.
- A CCJ replicates all the supported files inside the selected backup repository. If the intent is to replicate only the backup files from a specific job, then assign an exclusive backup repository to the job and create a CCJ that replicates that repository. To avoid creating too many Catalyst Stores, it is possible to create multiple Veeam Backup Repositories inside the same Catalyst Store. Replication occurs at the repository level and not at the entire Catalyst Store level.
- When the Veeam Backup job is writing to a replicated backup repository that has GFS retention, then the GFS retention files are replicated too.
- If you remove the entire source backup repository or delete all the restore points from the source as “Delete from disk,” then the replicated copies are not deleted.

Regarding parallel and cascade replication:

- To set up parallel replication from one backup repository to multiple destinations, it is necessary to set up an additional Catalyst Copy Job for each replication.
- To set up cascade replication from *Source* to *Copy 1* and from *Copy 1* to *Copy 2*, there are two equally viable options:
 - Option 1: Add a line to the same Catalyst Copy job in the **Storage Mapping** section.
 - Option 2: Set up an additional Catalyst Copy job from *Copy 1* to *Copy 2*.
- There is no predefined limit on the cascade sequence length; the only warning is to avoid loops, such as *Source* to *Copy 1* and *Copy 1* to *Source*.

Creating a Veeam-managed Catalyst Copy Job with HPE StoreOnce Cloud Bank Storage

When a Catalyst Copy Job (CCJ) is configured with HPE StoreOnce Cloud Bank Store as the backup target, Veeam initiates the Catalyst Copy operation from the local Catalyst Store to the specified Cloud Store backed by either on-premises or public object storage as depicted in Figure 42. As a Cloud Bank Store behaves similar to a Catalyst Store, the low cost and simplicity of object storage are further amplified by the HPE StoreOnce deduplication technology while sending, storing, and retrieving the data.




Figure 42. Catalyst Copy Job utilizing Cloud Bank Catalyst Store as the target repository

A HPE StoreOnce Cloud Bank Store needs to be created before configuring a CCJ with HPE Cloud Bank Store as the target repository. To create a Cloud Bank Store, log in to the HPE StoreOnce GUI and navigate to **Data Services → Cloud Bank Stores → Create Store**. In the **Service Provider Settings**, select a suitable object storage provider and populate the required fields to access the object storage, as displayed in Figure 43. Perform **Security Settings** and **Advanced Settings** as described in the [Creating a Catalyst store for Veeam backups](#) section.






Figure 43. Creating an HPE StoreOnce Cloud Bank Store

After successful creation of a Cloud Bank Store, it can be used for creating a target backup repository. See the [Creating an immutable HPE StoreOnce Catalyst-based Veeam backup repository](#) section for additional details. A Catalyst Copy Job for creating an off-site copy in HPE Cloud Bank Store can be done using the steps outlined in the [Creating a Veeam-managed Catalyst Copy Job with HPE StoreOnce Cloud Bank Storage](#) section.

An important feature of HPE Cloud Bank is that it is not dependent on the availability of the host (primary) HPE StoreOnce appliance. In the event an HPE StoreOnce appliance becomes unavailable or unusable, the backup copies contained in a detached or disconnected HPE Cloud Bank Store can be recovered using a separate HPE StoreOnce appliance, including an HPE StoreOnce VSA. This is possible because the local metadata held on the HPE StoreOnce appliance, which is primarily used for matching, is also duplicated across various objects in the cloud storage. A Cloud Bank Store can be re-attached to the same or a different HPE StoreOnce appliance using the following options:

- Connect Read Write: With this option, only one HPE StoreOnce appliance (physical/VSA) can access in RW mode. This is useful in case of site or hardware failures.
- Connect Read Only: With this option, multiple HPE StoreOnce appliances can concurrently access the same Cloud Bank Store in RO mode. This is particularly useful to test disaster recovery processes using a different HPE StoreOnce appliance using the backup data in the Cloud Bank Store. The access to data in RO mode does not have any impact—neither on the RW access from the original HPE StoreOnce appliance nor on the production services.

For more information on disaster recovery with HPE StoreOnce Cloud Bank Store, see [HPE Reference Configuration for HPE Cloud Bank Storage with S3 Connector and Scality RING](#).

Creating a Veeam Backup Copy Job with Catalyst writes over LAN/WAN to an HPE StoreOnce repository

Veeam Backup Copy Jobs are designed to copy backups of selected VMs or servers from any Veeam backup repository to a different Veeam backup repository. They can also set simple or tiered GFS retention. Configuring HPE StoreOnce as the target repository for Veeam Backup Copy jobs as depicted in [Figure 44](#) gives a whole set of unique advantages:



- **Bandwidth reduction over WAN**—Having the destination repository on HPE StoreOnce provides the advantage of using Catalyst source-side deduplication that greatly reduces the network bandwidth required for transferring data to a remote location. As shown in the Figure 44 configuration, the Veeam gateway server(s) running the Catalyst client should be located on the production site. The gateway server(s) reads backup data from the local repository, performs source-side deduplication, and writes to the remote HPE StoreOnce appliance. By doing so, deduplicated data is directly transferred over the WAN link.
- **Reduced storage space consumption**—Having the destination repository on HPE StoreOnce provides the advantage of the deduplication to reduce the storage utilization and related costs. The Backup Copy Jobs can be used to copy the backup from a short-term, fast-acting primary backup target to a highly deduplicated HPE StoreOnce secondary backup target.




Figure 44. Utilizing a Veeam gateway server to copy backups from a local HPE StoreOnce appliance to a remote HPE StoreOnce appliance

- **Two-tier configuration**—Veeam Backup Copy Jobs enable replicating the restore points from the primary backup repository to a second one in a two-tier deployment. The two-tier configuration shown in Figure 45 is ideal for environments where advanced restore features such as secure restore, IVMR, SureBackup and so on, are required. These functionalities are characterized by intensive random I/O traffic and can require a fast storage, such as HPE Apollo or HPE Alletra 5000 as the first-tier repository to provide more IOPS performance. In this configuration, the HPE StoreOnce target backup repository does not need to be located in a remote location. It can be co-located with the first-tier fast/short-term repository. Due to the performance advantages offered by the two-tier configuration, this backup solution design is considered a Veeam best practice.

Utilizing a Veeam snapshot-only backup job to keep short-term restore points in the form of storage snapshots on HPE Alletra 6000 and HPE Alletra 9000 reduces or eliminates the advantages of a two-tier configuration as the storage snapshots are even faster restore points than a first-tier backup repository and provide plenty of random I/O performance for multiple types of advanced restore and recovery operations.




Figure 45. Veeam best practice for high-frequency use of full-featured restore options with high random I/O characteristics

For additional information on two-tier data protection with HPE Nimble and HPE Apollo, see:

- HPE Nimble: [HPE Reference Configuration for Veeam Availability Suite with HPE Nimble Storage](#)
- HPE Apollo: [HPE Reference Architecture for Veeam Availability Suite with HPE Apollo backup target](#)



The configuration shown in [Figure 45](#) can either utilize a single Veeam gateway server or multiple gateway servers located on the production site. The gateway server(s) reads from the local repository and writes to the remote HPE StoreOnce appliance directly over the WAN. This can be achieved by either specifying **Automatic selection** or by excluding the servers that are not in the proximity of the source backup repository by utilizing the **Use the following gateway servers only** option as demonstrated in [Figure 46](#). When **Automatic selection** is configured, Veeam Backup & Replication version 12 automatically chooses the server that acts as the data mover for both the source and destination backup repositories as the gateway server.




Figure 46. Configuring gateway server for a Veeam backup repository

For the destination repository, enable **Gateway server and StoreOnce are connected over WAN**. For more information on creating an HPE StoreOnce Catalyst backup repository, see the [Creating an immutable HPE StoreOnce Catalyst-based Veeam backup repository](#) section. For more details on configuring a Veeam gateway server, see the [Deploying a Veeam gateway server](#) section.

After configuring the destination repository, the Backup Copy Job can now be configured. A BCJ is based on an “incremental forever” policy, and not on the usual forward incremental policy recommended for HPE StoreOnce. This policy is based on a single rolling full backup at the end of the chain.

The incremental forever has a few limitations:

- If incremental forever does not support multiple chains, then the daily retention is limited by the single-chain length supported by the target HPE StoreOnce appliance. This backup chain length varies from 21 to 42 restore points for an HPE StoreOnce Gen4+ model. See [Veeam Backup & Replication 12 - HPE StoreOnce](#) for more information.
- It is very important to set at least one weekly GFS. This is necessary as a “full backup file maintenance” to reclaim unused capacity on the rolling full backup. Without this step, the rolling full backup file will continue to grow forever, potentially using all the available capacity. Select the **Keep certain full backups longer for archival purposes** option and click **Configure...** as shown in [Figure 47](#) to define a GFS retention policy.

As shown in [Figure 47](#), the **Read the entire restore point from source backup instead of synthesizing it from increments** should not be selected, otherwise, the unused capacity reclamation may not work.



New Backup Copy Job

Target
Specify the target backup repository, number of recent restore points to keep, and the retention policy for full backups. You can use map backup functionality to seed backup files.

Job
Objects
Target
Data Transfer
Schedule
Summary

Backup repository:
ghn-corgi-bkprepo2 (Created by HPELAB\grishma at 11/7/2022 9:42 AM.)
146 TB free of 147 TB [Map backup](#)

Retention policy: 7 days

☒ Keep certain full backups longer for archival purposes
GFS retention policy is not configured
☐ Read the entire restore point from source backup instead of synthesizing it from increments

[Configure...](#)

Configure GFS

☒ Keep weekly full backups for: 1 weeks
Create weekly full on this day: Sunday

☐ Keep monthly full backups for: 1 months
Use weekly full backup from the following week of a month: First

☐ Keep yearly full backups for: 1 years
Use monthly full backup from the following month: January

[Advanced...](#)
[Cancel](#)

[Save as default](#) [OK](#) [Cancel](#)

Figure 47. Backup Copy Job “Target” details

The settings in the **Maintenance** tab shown in Figure 48 should be configured as follows:

Advanced Settings

Maintenance **Storage** RPO Monitor Notifications Scripts

Storage-level corruption guard
☐ Perform backup files health check (detects and auto-heals corruption) on:
At 22:00 on Last Saturday of every month [Configure...](#)

Full backup file maintenance
Use these settings to defragment and compact full backup file periodically when the job schedule does not include periodic fulls.
☐ Remove deleted items data after 30 days

[Save As Default](#) [OK](#) [Cancel](#)

Figure 48. Backup Copy Job advanced maintenance settings



- Do not enable **Storage-level corruption guard**. While technically supported, storage-level corruption guard requires reading the entire backup file back to the gateway server, which in this architecture is typically over a WAN where link speed and reliability might affect its performance.

Note

Catalyst checksum data is sent to prevent transmission errors. HPE StoreOnce has additional data integrity algorithms to continuously verify its storage to prevent silent corruptions. Also, RAID 6 prevents data loss, even after two concurrent disk failures.

- Do not enable **Full backup maintenance** for the same reasons mentioned above.

The settings in the **Storage** tab shown in Figure 49 should be configured as follows:

Advanced Settings

Maintenance | **Storage** | RPO Monitor | Notifications | Scripts

Data reduction

☐ Enable inline data deduplication (recommended)

Compression level:

Optimal (recommended) ▼

Optimal compression provides for best compression to performance ratio, and lowest backup proxy CPU usage.

Encryption

☐ Enable backup file encryption

Password:

▼ Add...

Manage passwords

Save As Default OK Cancel

Figure 49. Backup Copy Job advanced storage settings

- Do not check **Enable inline data deduplication** because deduplication is handled by HPE StoreOnce.
- Set the **Compression level** to **Optimal** and make sure that the backup repository has **Decompress backup data blocks before storing**. See the [Creating an immutable HPE StoreOnce Catalyst-based Veeam backup repository](#) section for more information.
- Do not select **Enable backup file encryption**. If the data on the source repository was already encrypted, the job will unencrypt before writing to the copy destination. If encryption for data-at-rest is required, use HPE StoreOnce encryption as described in the [Creating a Catalyst Store for Veeam backups](#) section.



On the BCJ **Data Transfer** section, to send data directly from the source repository to the target repository, select **Direct** as shown in Figure 50.




Figure 50. Specifying "Direct" as the "Data Transfer" type

Traditional Veeam Backup Copy Job configuration

The last option for copying remote backups involves the traditional Veeam process for backup replication, shown in Figure 51. This setup comprises of two Veeam Proxy servers: one at the local site and one at the remote site. With this scenario, Veeam uses its own bandwidth reduction features with optional WAN accelerators. Compared to the options from the previous two sections, this deployment requires significantly higher bandwidth and requires an additional server.




Figure 51. Traditional backup and replication setup with a Veeam proxy server at both the local and remote sites


This configuration can be achieved by configuring two different proxy/gateway servers:

1. **Source GW**—This server is in the source site, close to the source repository.
2. **Destination GW**—This server is in the destination site, close to the BCJ target repository on an HPE StoreOnce Catalyst Store.



For more details on configuring a Veeam Gateway server, see the [Deploying a Veeam gateway server](#) section. The Backup Copy Job configuration steps are the same as in the [Creating a Veeam Backup Copy Job with Catalyst writes over LAN/WAN to an HPE StoreOnce repository](#) section.

When the available bandwidth is very limited, it is also possible to configure the optional Veeam **Source WAN accelerator** and a **Target WAN accelerator**, as shown in Figure 52. For more details on WAN accelerator configuration, see [Veeam Backup & Replication 12 - WAN Accelerators](#).



Edit Backup Copy Job Backup Copy Job 1

Data Transfer
Choose how object data should be transferred from source to target backup repository.

Job
Objects
Target
Data Transfer
Schedule
Summary

☐ **Direct**
Object data will be sent directly from source to target repository. This mode is recommended for copying backups on-site, and off-site over a fast connection.

☒ **Through built-in WAN accelerators**
Object data will be sent to target repository through WAN accelerators that must be deployed in both source and target sites. This mode provides for significant bandwidth savings.

Source WAN accelerator:
[Dropdown menu]

Target WAN accelerator:
[Dropdown menu]

< Previous Next > Finish Cancel

Figure 52. Veeam Backup Copy Job that uses WAN accelerators to transfer object data to the target backup repository

Data transfer between the two sites is not highly deduplicated as in the previous solution designs based on HPE StoreOnce Catalyst Copy. For this reason, when applicable, the above use cases are preferable.

Recovering data from an HPE StoreOnce Catalyst repository

Veeam Backup & Replication handles backup files saved on backup repositories based on HPE StoreOnce Catalyst Stores the same as repositories based on other storage types, such as local disks or NAS. HPE StoreOnce Catalyst Store backups can be used for Instant VM Recovery (IVMR), Veeam On-Demand Sandbox (including SureBackup), restoring guest VM files, restoring the entire VM, restoring virtual disks, and other restore options offered by Veeam Backup & Replication. In the event of a malicious ransomware attack, it is recommended to use the Veeam **Secure Restore** feature using **immutable HPE StoreOnce backup copies** to make sure that no threats are re-introduced into the production systems after a successful recovery. For more information regarding restore performance, see the [Guidance of the use of HPE StoreOnce Systems with Veeam Backup & Replication technical white paper](#).

Migrating Veeam backup repositories to/from an HPE StoreOnce Catalyst Store

There is an easy and effective methodology for migrating entire Veeam backup repositories to new storage platforms and even across different storage types. The migration process takes care of migrating all the backup files and the associated restore points. More information can be found at: [Veeam Backup & Replication 12 - Evacuating Backups from Performance Extents](#).

Examples of supported migrations include:

- Existing repository based on NAS/NTFS/XFS/ReFS → new HPE StoreOnce Catalyst-based backup repository
- Existing HPE StoreOnce Catalyst-based backup repository → new repository based on NAS/NTFS/ReFS
- Existing HPE StoreOnce Catalyst-based backup repository → new HPE StoreOnce Catalyst-based backup repository



Some limitations of this methodology include:

- The process migrates entire backup repositories. It is not possible to select data subsets.
- The process does not modify the data format. The existing data format must be compatible with the target repository. Specifically, compression, deduplication, encryption, backup policy, and chain length are not changed. Before starting the migration to HPE StoreOnce, be sure that the source format is compatible with the target.
- Make sure that the format of each Veeam backup job is compatible with the settings outlined in [Creating a Veeam backup job for VMs](#).

The migration process is managed by Veeam Backup & Replication and requires the steps summarized below.

1. Create or use a Scale Out Backup Repository (SOBR),
2. Add an HPE StoreOnce Catalyst repository as a new extent. This will be used as the destination for migration.
3. Put the old extent in Maintenance mode and evacuate it.

The following example demonstrates how to migrate a repository to an HPE StoreOnce Catalyst-based repository from a file system-based repository:

1. Navigate to **BACKUP INFRASTRUCTURE**, right-click **Scale-out Repositories**, and select **Add scale-out backup repository...**, demonstrated in Figure 53.




Figure 53. Adding a scale-out backup repository

2. Provide a **Name** and optional **Description** to the scale-out backup repository, illustrated in Figure 54.





New Scale-out Backup Repository

Type in a name and description for this scale-out backup repository.

Name

Performance Tier

Placement Policy

Capacity Tier

Summary

Name:

Disk-to-Catalyst

Description:

Created by HPELAB\grishma at 11/21/2022 3:13 AM.

< Previous

Next >

Finish

Cancel

Figure 54. Providing a name and description for the scale-out backup repository

3. Add the source repository and the target repository, shown in Figure 55. The order that the repositories are added does not matter.

[illegible]


Figure 55. Selecting the "Performance Tier" for the scale-out backup repository



Note

After the backup repository is added to the scale-out backup repository, Veeam automatically updates backup jobs utilizing the backup repositories added in [Figure 55](#) to this newly created scale-out backup repository.

- Make sure to set **Use per-machine backup files** on the advanced settings of Figure 56.



Advanced Settings [X]


☒ **Use per-machine backup files (recommended)**
Improves backup performance for storage devices benefiting from multiple I/O streams. This is the recommended setting when backing up to enterprise grade block storage and deduplicating storage appliances.

☐ **Perform full backup when the required extent is offline**
When a repository extent with previous backup files cannot be reached, jobs will perform an active full backup instead of failing to create an incremental backup. Scale-out repository must be sized appropriately to handle multiple fulls.

[OK] [Cancel]

Figure 56. Advanced performance tier settings of a Veeam scale-out repository

- Select **Data locality** to migrate all backup files to the target repository, demonstrated in Figure 57.



New Scale-out Backup Repository [X]

Placement Policy
Choose a backup files placement policy for this performance tier. When more than one extent matches the placement policy, backup job will choose the extent with the most free disk space available.

Name

Performance Tier

Placement Policy

Capacity Tier

Summary

☒ **Data locality**
All dependent backup files are placed on the same extent. For example, incremental backup files will be stored together with the corresponding full backup file. However, the next full backup file can be created on another extent (except extents backed by a deduplicating storage).

☐ **Performance**
Incremental backup files are placed on a different extent from the corresponding full backup file, providing for better backup file transformation performance with raw storage devices. Note that losing an extent with a full backup makes restoring from increments impossible.
Specify the placement policy for full and incremental backup files. [Customize...]


☐ **Strict placement policy enforcement**
By default, whenever we are unable to place a new backup according to the policy, we will violate it to ensure a backup could still be created. Select this option if you'd rather have backup jobs fail in such circumstances instead.

[< Previous] [Next >] [Finish] [Cancel]

Figure 57. Choosing a "Placement Policy" for the scale-out backup repository

- For the **Capacity Tier** shown in Figure 58, do not select **Extend scale-out backup repository capacity with object storage**.





New Scale-out Backup Repository

Capacity Tier
Specify object storage to copy backups to for redundancy and DR purposes. Older backups can be moved to object storage completely to reduce long-term retention costs while preserving the ability to restore directly from offloaded backups.

Name

Performance Tier

Placement Policy

Capacity Tier

Summary

☐ **Extend scale-out backup repository capacity with object storage:**
 <Click Choose to pick object storage> **Choose...**

Define time windows when uploading to capacity tier is allowed **Window...**

☐ **Copy backups to object storage as soon as they are created**
 Create additional copy of your backups for added redundancy by having all backups copied to the capacity tier as soon as they are created on the performance tier.


☒ **Move backups to object storage as they age out of the operational restore window**
 Reduce your long-term retention costs by moving older backups to object storage completely while preserving the ability to restore directly from offloaded backups.
 Move backup files older than **14** days (your operational restore window) **Override...**

☐ **Encrypt data uploaded to object storage**
 Password: **Add...**
Manage passwords

< Previous **Apply** **Finish** **Cancel**

Figure 58. Optionally enabling object storage with a scale-out repository

- Review the summary and click **Finish**.



New Scale-out Backup Repository

Summary
Review the scale-out backup repository settings, and click Finish to exit the wizard.

Name

Performance Tier

Placement Policy

Capacity Tier

Summary

Summary:
 Scale-out backup repository was created successfully.

< Previous **Next >** **Finish** **Cancel**

Figure 59. The “Summary” page of the scale-out backup repository wizard



- Referring to Figure 60, click the newly created scale-out repository containing the source repository and the target repository. In this scenario, “ghn-corgi-bkprepo1” was the target repository, which was an HPE StoreOnce Catalyst Store, and “Backup Repository 1” was the source repository, which was disk storage on a Windows server.




Figure 60. Selecting a scale-out repository

- Right-click the source repository and put the source repository into **Maintenance mode** as shown in Figure 61. This is necessary to avoid the preconfigured backup jobs use this extent as the backup target.




Figure 61. Placing a backup (source) repository into maintenance mode



10. Once in **Maintenance mode**, right-click the source repository and select **Evacuate backups**, demonstrated in Figure 62. The operation starts to migrate the backup files from the source repository to the destination repository.




Figure 62. Evacuating the backups from a repository into a different Veeam repository

11. As shown in Figure 63, Veeam will now migrate the backups from the source repository to the target repository.




Figure 63. Veeam executing backup evacuation and migrating the backups to a separate repository



12. After the evacuation process is complete, it is possible to edit the scale-out backup repository and **Remove** the source extent, as shown in Figure 64.




Figure 64. Removing the source repository from a scale-out backup repository

After the migration is completed, there is no need to convert the destination scale-out backup repository to a normal repository. The backup jobs that previously wrote to the migrated repository will now use the destination repository. All the restore points are visible and available for restore operations, and the process for deleting expired restore points should work as expected.

Note

The migration process can run quite fast and use a considerable amount of system resources, especially when there is a large amount of data to migrate. If the goal is not to complete the migration as fast as possible, but to avoid a slow down in other production activities, then it is possible to set temporary limits in the **Load control** section as demonstrated in [Figure 65](#). Remember to remove the limits after the migration is complete.



Edit Backup Repository

Repository
Type in path to the folder where backup files should be stored, and set repository load control options.

Name
HPE StoreOnce

Repository

Mount Server

Review

Apply

Summary

Location
Catalyst store:
CSDeflmmutability

Capacity: <Unknown> Free space: <Unknown> **Populate**

☒ Make recent backups immutable for: 7 days
Protects backups from modification or deletion by ransomware, malicious insiders and hackers. GFS backups are made immutable for the entire duration of their retention policy.

Load control
Running too many concurrent tasks against the repository may reduce overall performance, and cause I/O timeouts. Control storage device saturation with the following settings:

☒ Limit maximum concurrent tasks to: 4

☐ Limit read and write data rate to: 1 MB/s

Click Advanced to customize repository settings. **Advanced...**

< Previous **Next >** Finish Cancel

Figure 65. Limiting the maximum concurrent tasks for a Veeam repository

Configuring Veeam Tape Jobs with HPE StoreEver

The HPE StoreEver tape libraries keep data secure with hardware-based data encryption, WORM media, and removable storage that can be moved offline to avoid threats such as cyber-attacks or site disasters. Tapes—especially when they are ejected from the library and stored in a vault—are the best example of air-gapping. Modern data protection solutions can include tapes to comply with all the conditions in the 3-2-1-1 rule, especially for the last “1” requiring an off-line or air-gapped backup copy. Tapes are by far the most cost-effective storage for data protection, and they can be used for creating an effective storage tier for long and very long-term retention backup.

Tapes media is optimized for sequential read/write access and is very fast. For example, LTO-9 has an improved hardware compression technology over the previous generations. This helps to increase the extent of data reduction and backup speed for data not pre-compressed by Veeam. On Veeam workloads, an LTO-9 has the following throughput characteristics:

- **300 MB/s as maximum native throughput**

This is the maximum throughput that can be achieved when a tape job copies already pre-compressed backup data or when the hardware compression on the tape drive is disabled on Veeam Backup & Replication tape jobs.

- **1000 MB/s as maximum throughput**

This is the maximum throughput that can be achieved only when the backup data is highly compressible by the tape drive.

Sometimes a tape drive can sustain a throughput that exceeds the capabilities of the surrounding infrastructure, and the tape device has to stop, rewind, and restart several times during the read/write process. This leads to a problem called “shoe shining.” To avoid this, an HPE StoreEver LTO-9 Ultrium tape drive is equipped with Data Rate Matching. This feature combines advanced buffer management with HPE exclusive Variable Tape Speed technology to further optimize performance by matching the speed of source systems to keep drives streaming anywhere between 100 MB/s and 300 MB/s (native data). Without this feature, when the infrastructure surrounding the Veeam Backup & Replication Tape Server—including primary storage, backup repositories and LAN—cannot sustain the drive data stream throughput, the overall performance could degrade significantly. The continuous start, stop, and rewind process could also increase the wear on the tape media and tape device, reducing its life and reliability. These features make HPE StoreEver tape libraries a fast, highly reliable, and cost-effective target for Veeam backup solutions.

Veeam Backup & Replication users can copy/write backup files of VMs and Veeam agents stored in backup repositories and regular files in Windows and Linux servers to tape for long-term archival or disaster recovery purposes. HPE StoreEver has full-fledged integration with Veeam as an offline target for the golden copies created through:

- **Backup to tape job**—Jobs for copying backup files from a disk-based storage, such as HPE StoreOnce or HPE Apollo 4000, to tape.
- **Files to tape job**—Jobs for copying regular files in Windows and Linux servers to tape for long-term archival or disaster recovery purposes.

Before configuring end-to-end data protection and an archiving solution with HPE StoreEver, see the [HPE StoreEver Compatibility Matrix](#).

Adding a Tape Server

A Tape Server interconnects a tape device to the Veeam Backup & Replication server. It hosts the Veeam data-mover component that enables the transfer of data during archival and restore operations from the tape media. Veeam automatically discovers and catalogs the HPE StoreEver tape library and tape media after the Tape Server is added to the infrastructure.

By default, the role of the Tape Server is assigned to the backup server having FC or SAS connectivity to the tape device. However, in the case of large and complex environments, it is recommended to use a dedicated Tape Server with partitions configured in the tape library.

To add a Tape Server, go to the **Tape Infrastructure** view, right-click **Tape Servers** and select **Add tape server...** as displayed in Figure 66. Populate required details and select **Start tape libraries inventory when I click Finish** to enable auto-discovery of attached tape devices and media. For more details, see [Veeam Backup & Replication 12 – Adding Tape Servers](#).




Figure 66. Adding tape server and auto-discovery of tape devices

After the library is auto-discovered, it is populated under the **Libraries** node as shown in Figure 67.






Figure 67. Listing of auto-discovered HPE StoreEver tape libraries under library node

Note
Make sure that the appropriate HPE StoreEver native tape device driver is installed on the Tape Server before adding it to the Veeam Backup & Replication server. See [HPE StoreEver Tape Drivers](#) for more details.

Creating a custom Media Pool

After the auto-inventory of the tape libraries and tape media is completed successfully, or when new tapes are loaded into the tape library, they are presented as Unrecognized. Use the tape inventory feature to identify unrecognized tapes. After inventory, tapes are displayed in a free Media Pool. For improved management and tracking of tapes, a custom Media Pool can be created by utilizing the tapes in the free pool:

1. Click the **Media Pools** node, then select **Add media pool** to launch the **New Media Pool** wizard to create a standard Media Pool or WORM Media Pool using HPE StoreEver WORM tapes.
2. Add desired tapes from the free Media Pool in the **Tapes** tab. Select **Add tapes from the free Media Pool automatically when more tapes are required**, as shown in Figure 68.






Figure 68. Adding free tapes to custom Media Pool

- 3. Name the Media Set, select the frequency for creation of Media Set, and set a Retention Policy according to administrator requirements.
- 4. Enable parallel processing for the Media Pool as displayed in Figure 69 so that multiple Tape jobs or multiple backup chains within a single job are processed by multiple tape drives in the selected libraries.

Note
A Veeam Backup & Replication tape job feeds each tape with a single backup file at a time. This means that the data source can limit the tape job throughput. Essentially, a tape cannot write/backup faster than the source repository. The same consideration is valid for restore processes, where the entire pipeline to the destination storage needs to sustain the tape throughput.

Additionally, select the **Use encryption** checkbox to enable hardware encryption of data stored in the tape. With this setting, tape data is further secured by the AES 256-bit hardware encryption performed by the HPE StoreEver LTO-9 Ultrium tape drive. Enabling this option does not reduce tape job performance and eliminates the need for separate hardware or software.

When the source repository is an HPE StoreOnce Catalyst Store, the backup data is usually not pre-encrypted by Veeam, as this adversely impacts deduplication. Even when Catalyst Store encryption is enabled, data is decrypted as it is retrieved from the store. For this reason, enabling the tape drive hardware encryption is highly recommended to avoid the situation where a misplaced tape media can be used for reading the confidential data it might contain.



Edit Media Pool LTO9Pool ✕

Options
Specify parallel processing and encryption settings.

Name

Tapes

Media Set

Retention

Options

Summary

☒ Enable parallel processing for tape jobs using this media pool

Jobs pointed to this media pool can use up to tape drives simultaneously

☒ Enable parallel processing of backup chains within a single tape job

Backup to Tape job populated with multiple source backup jobs or with a job using per-VM backup file chains will process separate chains in parallel using one media set per drive.

☒ Use encryption

Password:

Created by HPELAB\grishma at 6/17/2022 7:28 AM. (Last edited: 52 days ago) Add...

Loss protection disabled Manage passwords

Hardware encryption will be used if tape drive supports it. Otherwise, software encryption will be used, which may result in small performance penalty.

< Previous
Apply
Finish
Cancel

Figure 69. Enabling parallel processing and encryption for the Media Pool

Note

Veeam Backup & Replication allows the user to create a GFS Media Pool by using GFS tape jobs to archive data to tape according to a GFS backup rotation scheme. For more details on GFS Media Pools, see [Veeam Backup & Replication 12 - GFS Media Pools](#).

Creating a backup to tape job

Veeam Backup & Replication software cannot directly backup VMs, physical machines, and Veeam agents to tape devices. A Veeam backup repository that is the target of a Backup Job, or a Backup Copy Job, or a Catalyst Copy Job, or the jobs themselves can be used as a source for archiving data to tape. Some of the common solution designs are shown below.

In the topology shown in Figure 70, an HPE StoreOnce Catalyst repository or primary backup job created with HPE StoreOnce as the backup target can be used as a source for backup to tape jobs. The tape job can also be linked to a primary backup job to automate the backup and archive workflow.




Figure 70. Archiving backup data from HPE StoreOnce to HPE StoreEver

In a multisite configuration as shown in Figure 71, HPE StoreEver can be used as a centralized repository for air-gapped LTR of backups. In this case either the HPE StoreOnce repository that is used as the target for multiple Catalyst Copy Jobs or the Catalyst Copy Jobs can be specified as the source for the Backup to Tape jobs. This solution design offers multiple advantages, including cost and complexity reduction, because the tape library is in the central location only.




Figure 71. Multi-site backup consolidation using HPE StoreOnce and HPE StoreEver

To configure a new backup to tape job:

1. From the **Home** tab, click **Tape Job** on the top ribbon and select **Backups**.
2. Provide a suitable **Name** and **Description**, in the **Backups** tab click **Add...**, then select either **Backup jobs...** or **Backup repositories...** as a source for the job, as shown in Figure 72.




Figure 72. Specifying the source for machine backup to tape jobs

3. Specify a Media Pool for full and incremental backups.
4. In the **Options** tab, enable the settings as displayed in Figure 73.
 - a. **Eject tape media upon job completion:** Selecting this option automates the placement of a written tape into a free tape device slot.
 - b. **Export current media set upon job completion:** This option needs to be used if the tape media with golden copies of data need to be physically moved to an air-gapped storage location. It is possible to schedule the export of the media set to mail slots in the HPE StoreEver tape library on desired days as well.



New Backup to Tape Job

Options
Specify tape job options.

Backups
Media Pool
Incremental Backup
Options
Schedule
Summary

Media automation

☒ Eject tape media upon job completion
This option makes the job automatically eject tape from drives upon completion, so that tape does not stay in the drive, which is a best practice.

☒ Export current media set upon job completion Days...
This option makes the job automatically close and export the current media set on specific days.

Drive usage

☒ Limit the number of drives this job can use to: 2 tape drives

Advanced job settings include compression, notification settings, automated post-job activity and other settings. Advanced...

< Previous **Next >** Finish Cancel

Figure 73. Enabling options for automated air-gapping of tape media

5. In the **Advanced Settings** in the **Options** tab, it is possible to select **Process latest full backup chain only** as displayed in Figure 74 to only the latest backup chain, thereby minimizing space consumption of the tape media. Hewlett Packard Enterprise recommends enabling **Use hardware compression when available** because it expedites the archival process and minimizes the space utilization on tape media.






Figure 74. Options for optimizing tape space consumption and improving archiving speed

Note

For further optimization of duration and space required for the archival process, schedule the machine backup to tape job right after the periodic synthetic full/active full backup creation and before the next incremental backup schedule is triggered. By doing so, the Veeam tape job copies only the latest full backup file instead of copying the entire latest backup chain, thereby reducing the time and space required to copy files from the primary backup repository to the HPE StoreEver tape library.

Creating a file to tape job

A Veeam Backup & Replication file to tape job allows the administrator to directly backup specified files and directories from a Veeam-managed server—Windows or Linux—to tape media. In the case of file to tape jobs, data is directly transferred via LAN to the Tape Server and then to the SAN or SAS-connected tape media. To create a file to tape job, perform the following steps:

1. From the **Home** tab, click **Tape Job** on the top ribbon and select **Files**.
2. Provide a suitable **Name** and **Description**. In the **Files and Folders** tab, select the required files or folders that need to be backed up.
3. In the **Full Backup** and **Incremental Backup** tabs, specify the desired target tape media pool and schedule for full and incremental backups.
4. In the **Options** tab, select **Eject tape media upon job completion** and **Export current media set upon job completion on selected days**. These options aid in air-gapping the media sets containing the backup data.

For more details, see [Veeam Backup & Replication - File Backup to Tape](#).

Recovering data from HPE StoreEver tape media

The machine or VM backup data can be restored from the tape media using any of the following options:

- Restore directly to infrastructure—This option is recommended only when a small-sized VM or a small number of VMs need to be recovered from a backup consisting of data from several VMs.



- **Restore through a staging repository**—Accessing data randomly from a disk-based repository is much quicker than accessing data directly from tape. This option is preferred when a large number of VMs, physical machines, or an entire Veeam agent needs to be recovered. Recovery of data using this option is a two-step process: in the first step, data is copied from the tape to a staging repository or a folder on disk, and in the subsequent step, Veeam initiates a restore of the entire VMs or desired files to infrastructure.
- **Backup restore from tape to repository**—This option allows you to copy VM backups from tape to repository. This is helpful if you need some backups on disk for later use, or also for VM guest OS files restore. The restored backup is populated under the imported disk backup node on the Veeam Backup & Replication console.

While performing restores from tape through a staging repository, the HPE StoreOnce Catalyst Store can be selected as the intermediate backup repository. In cases where you want to use the restore points for activities that generate an intense random I/O workload, such as IVMR or complex granular restore operations, it might be better to stage the restore to a normal file system on a fast storage. This way you avoid creating excessive resource utilization on HPE StoreOnce, which could slow down other concurrent workloads.

Restoring files from tape backups is like a traditional backup application in that you choose a file for restore and destination. That file is then restored to the destination directly from tape. As with the case of backing up files to tape, the file movement first occurs over the LAN to the Veeam backup server. Long restore times are possible because of LAN bottlenecks.

For more details on the procedure to recover machine backup data and files from tape, see [Veeam Backup & Replication 12 - Tape Devices Support](#).

Performance and tuning considerations

The entire backup infrastructure needs to be analyzed and tweaked for maximizing the performance of the archival process to HPE StoreEver tape libraries. Some of the crucial factors that determine the maximum throughput of a tape job include:

- **Configuration of backup source or restore destination storage:** With the evolution of LTO technology, tape devices are capable of very high throughput on sequential workloads. Data cannot be put onto tape any faster than it can be read from the source disk. When an HPE StoreOnce appliance is used as the primary deduplicating backup target, it is recommended to use a 25GbE LAN connection between the HPE StoreOnce appliance and the gateway server. Hosting the Catalyst gateway services on the same server as the tape server avoids an additional hop in LAN and this optimization improves the tape job throughput.
- **Connectivity to the tape library:** For complete SAN integration, it is recommended to use a Fibre Channel connection between the tape server and HPE StoreEver tape library with a dedicated partition.
- **Compressibility of the data:** The compression ratio determines the amount of data that can be stored on each tape cartridge and how efficiently tape drives can read and write records to the media. The higher the compression ratio, the higher the data transfer rate to and from the tape drive.
- **Average file size and fragmentation:** For file to tape jobs, backup and restore throughput could decrease when there is large number of small files, increased directory depth, and a high level of file fragmentation on the source/destination file systems.
- **Generation of the tape drive:** HPE StoreEver LTO tape drives have varying levels of performance depending on the generation. For more details, see [HPE StoreEver LTO Ultrium Tape Drive](#).

Note

Often the source systems tend to be the bottleneck for realizing the maximum speed of tape drives. To understand more about performance tuning, see [HPE Library and Tape tools](#).

Performance examples

To help you properly size a backup to tape solution for your requirements, a few backup tests have been performed and measured to evaluate the performance on different use cases. The measured throughput should not be considered the maximum possible, but what has been measured in the HPE lab with the configuration described below. During the tests, there were no other jobs running to create resource contention.

The HPE StoreOnce system used for this test has only one disk expansion with 11 active disks, excluding spares. HPE StoreOnce read performance is generally higher on a system with more disks. A small configuration was intentionally tested to make sure that even in this case, the overall throughput was enough to keep the LTO-9 in streaming mode. A system with more disks could show higher throughputs.



The tape server also hosted the Veeam Catalyst gateway services with 25GbE connectivity to HPE StoreOnce 5260 and 16G FC connectivity to the HPE StoreEver MSL3040 tape library.

Table 2 shows the Veeam backup to tape job throughput as shown on the Veeam GUI. Values are in MiB/s.

Table 2. Performance (write throughput) of Veeam Backup & Replication tape jobs with HPE StoreEver

Type of Tape Job	Source system configuration	Data format on source repository	Compression ratio of backup data in source repository	Is tape compression flag active?	Average logical write throughput as reported by CVTL (MB/s)	Veeam Tape Job peak processing rate (MiB/s)	Veeam reported bottleneck
Machine backup to tape job (Single VM full backup to tape)	HPE StoreOnce 5260 with 1 disk expansion and 25GbE connectivity to tape server	Data not pre-compressed and not encrypted	2.0:1	Yes	370	430	Source
Machine backup to tape job (Single VM full backup to tape)	ReFS Backup repository on HPE ProLiant DL380 backed by RAID 6 NVMe virtual volume and 25GbE connectivity to tape server	Data not pre-compressed and not encrypted	2.0:1	Yes	470	570	Proxy
Machine backup to tape job (Single VM full backup)	repository on HPE ProLiant DL380 backed by RAID 6 NVMe virtual volume	Data compressed and encrypted	2.0:1	Yes	270	270	Target
Machine backup to tape job (2 VMs parallel full backup to tape)	HPE StoreOnce 5260 with 1 disk expansion and 25GbE connectivity to tape server (same as the Catalyst Gateway server)	Data not pre-compressed and not encrypted	2.0:1	Yes	546 (Drive1 - 283, Drive2 - 263)	643	Source
File to tape job (direct backup to tape)	NTFS disk on HPE ProLiant DL380 backed by RAID 6 NVMe virtual volume	N/A	1.9:1	Yes	550	760	Target

Summary

Veeam Backup & Replication provides a solid software package for backup, replication, and recovery; however, without effective storage infrastructure, meeting the demands of the business can be difficult. The HPE StoreOnce purpose-built backup appliance (or VSA) is a powerful storage solution with a carefully designed end-to-end functional integration with Veeam Backup & Replication to provide a wealth of benefits to an organization of any size or complexity. HPE StoreOnce is a mature solution that offers unbeatable reliability, huge consolidation, and highly simplified management. The integration process with HPE StoreOnce and Veeam Backup & Replication is based on co-developed code and APIs. Starting from Veeam Backup & Replication Version 9, with every new release, the integration has offered new exclusive functionalities.

HPE StoreOnce offers the highest possible deduplication and a lower cost per TB for Veeam workloads, along with a set of powerful integrated features. For example, Catalyst Copy is an HPE StoreOnce feature that seamlessly integrates with Veeam Backup & Replication software to replicate the deduplicated backups to multiple remote Catalyst Stores for off-site protection without burdening any of the backup infrastructure components. Veeam Backup & Replication Catalyst Copy jobs further extend its support to HPE StoreOnce Cloud Bank Storage, enabling customers to replicate backup to cloud object storage, maintaining the same deduplication level as that of the local HPE StoreOnce repositories. Customers can create Disaster Recovery and Long-Term Retention solutions in the cloud, and thanks to HPE StoreOnce deduplication, you need to provision only a fraction of the cloud storage capacity required by other solutions such as Veeam's own cloud tiering.

The most salient new feature of HPE StoreOnce integration with Veeam Backup & Replication Version 12 is its ability to completely isolate and protect data from ransomware attacks through powerful role-based access control, dual authorization, and strong object-level immutability. HPE StoreOnce Catalyst APIs are leveraged by Veeam Backup & Replication to create immutable backups that give confidence no hacker—including the ones who have access to the credentials of all the Administrator accounts—can neither delete nor corrupt the backup data. The backup data remains intact for a successful recovery process in the event of a ransomware attack.



HPE StoreEver storage with Veeam Backup & Replication can provide a completely automated backup to tape solution. HPE StoreEver tape systems offer significant cost, energy, and footprint advantages while addressing data growth sustainably. HPE StoreEver tape systems integrate easily with Veeam Backup & Replication to provide measurable high performance and a great set of options, along with scalability from small to huge requirements. With HPE StoreEver and Veeam, customers can quickly store backup for air-gap protection and for long-term retention using the latest LTO tape technologies and taking advantage of the lowest cost per TB versus other on-premises or cloud-based solutions.

With HPE StoreOnce and HPE StoreEver storage systems, Hewlett Packard Enterprise, together with Veeam, offers end-to-end data protection solutions that are carefully integrated, optimized, and tested by two leading companies, and appreciated year after year by thousands of customers all over the world.



Resources, contacts, or additional links

HPE StoreOnce Data Protection Backup Appliances, hpe.com/storage/StoreOnce

HPE StoreOnce Gen4+ Systems User Guide,
support.hpe.com/hpesc/public/docDisplay?docId=sd00001027en_us&page=getting_started_chapter.html

HPE StoreOnce Compatibility Matrix, hpe.com/Storage/StoreOnceSupportMatrix

HPE StoreEver Compatibility Matrix, h20272.www2.hpe.com/SPOCK/Pages/spock2Html.aspx?htmlFile=hw_storeever.html

HPE StoreEver tape storage products, hpe.com/us/en/storage/storeever-tape-storage.html

Protecting Data from Ransomware with HPE StoreOnce Catalyst, hpe.com/V2/getpdf.aspx/A00042003ENW.pdf

HPE StoreOnce (Veeam Help Center), helpcenter.veeam.com/docs/backup/vsphere/deduplicating_appliance_storeonce.html

HPE Storage Solutions with Veeam, veeam.com/hpe-storage-solutions.html

Long-Term Retention Policy (Veeam Help Center), helpcenter.veeam.com/docs/backup/vsphere/gfs_retention_policy.html

HPE Reference Configuration for HPE 3PAR Storage and HPE StoreOnce with Veeam Availability Suite,
hpe.com/v2/getdocument.aspx?docname=a00061747enw

HPE Reference Configuration for Veeam Availability Suite with HPE Nimble Storage,
hpe.com/v2/getdocument.aspx?docname=a00079582enw

HPE Storage, hpe.com/storage



HPE Networking, hpe.com/networking

HPE Technology Consulting Services, hpe.com/us/en/services/consulting.html

Learn more at

hpe.com/solutions

**Make the right purchase decision.
Contact our presales specialists.**



© Copyright 2023 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Oracle is a registered trademark of Oracle and/or its affiliates. SAP and SAP HANA are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. Veeam Backup & Replication, Instant VM Recovery, On-Demand Sandbox, and SureBackup are registered trademarks or trademarks of Veeam. VMware and VMware ESXi are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. Windows and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. All third-party marks are property of their respective owners.